
datapro
**NETWORK
MANAGEMENT**

**VICE PRESIDENT-
EDITOR IN CHIEF**
J. Richard Peck

Product Manager
Algis V. Salciunas

Group Managing Editor
Lester P. SIZES

Associate Managing Editor
Lance Lindstrom

Associate Editor/Analyst
Jill Ann Huntington

Editorial Secretary
Phyllis Jack

Editorial Publications Coordinator
Lorraine A. Reese

PRESIDENT & PUBLISHER
Bruce R. Hollows

Vice President, Sales & Marketing
Martin J. Murphy

Vice President, Planning & Development
Donald F. Welsher

Vice President, International
Carl G. Tobiasen

Director, Asian Operations
A. Rajendra

Director, Canadian Operations
Steven T. Webb

Director, European Operations
Laurence Blackall

CUSTOMER SERVICE

Telephone: (800) 328-2776
Manager, Customer Service
Kathleen C. Patto

Canada (416) 298-1177
Switzerland (41) 21 27 41 71
United Kingdom (44) 628 773 277

DATAPRO RESEARCH □ 1805 Underwood Blvd. □ Delran, NJ 08075 USA □ (800) 328-2776 □ Telex 4761231 DPRO UI.

AUSTRALIA: McGraw-Hill Book Co., 4 Barcoo St., Roseville East, NSW 2069, Australia. Telephone (61) 2 406-4288. Telex 20849. MCGRAW AA.

CANADA: Datapro Research, 330 Progress Avenue, Scarborough, Ontario M1P 2Z5. Telephone (800) 668-9308.
Telex 06525169 MCGRAWHILL TOR.

JAPAN: Nikkei Business Publications Subscription Sales Co., 1-14-6 Uchikanda, Chiyoda-ku, Tokyo 101, Japan. Telephone (81) 3 233-8081. Telex 29902 NIKKEIBPJ.

SINGAPORE: Datapro Research, Unit 05-03 Dapenso Building, 158 Cecil Street, Singapore 0106. Telephone (65) 222-5091. Telex 21912 DPRO RS.

SWITZERLAND: Datapro Research, SA, Case Postale 460, CH-1000 Lausanne 17, Switzerland. Telephone (41) 21 27 41 71. Telex 458196 DSSA CH.

UNITED KINGDOM: Datapro Research, McGraw-Hill House, Shoppenhangers Road, Maidenhead, Berkshire SL6 2QL England. Telephone (44) 628 773 277. Telex 848484 MCHILL G.

datapro®

COPYRIGHT © 1989 McGRAW-HILL, INCORPORATED
DATAPRO RESEARCH, DELRAN, NJ 08075 USA
REPRODUCTION PROHIBITED



DATAPRO NETWORK MANAGEMENT.

COPYRIGHT © 1989 by McGRAW-HILL, INCORPORATED. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base retrieval system, without the prior written permission of the publisher.

Information has been obtained by Datapro Research/Datapro Network Management from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Datapro Research, or others, Datapro Research does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or for the results obtained from use of such information.

Meet Our Advisory Board

We are proud to present our Advisory Board for *Datapro Network Management*. Each of our advisors has made a commitment to a continuing involvement in the ongoing publication of bi-monthly issues. Our advisors are an active resource to us, providing ongoing review of the contents of all issues, informal advice about the technical and editorial direction of the service, and, in many cases, contributing reports for publication. Many thanks to each of them for their continuing support.

SANJIV BHATNAGAR

Mr. Sanjiv Bhatnagar is Senior Partner of International Systems Integration (ISI) Inc., a professional services firm based in Redmond, WA. The firm provides consulting and custom development in the areas of network management, vendor interconnectivity, and integration of information and control systems. The firm has performed work for clients, analyzing their network management systems and defining architectures and product strategies. Prior to joining ISI, Mr. Bhatnagar was involved in the development of operating system internals for the Meridian SL-1 voice messaging system while at Bell Northern Research (BNR) and the design of a multivendor Network Management System at Network Equipment Technologies (N.E.T.). He received his master's degree in computer science from Colorado State University.

CHRIS COLE

Chris Cole is president and cofounder of Peregrine Systems, Inc., located in the greater Los Angeles area. The company initially developed network management products as well as operating systems, Basic interpreters, database management systems, report writers, and application software for the IBM Series 1 computer. Mr. Cole developed the first commercially available symbolic manipulation language (SMP) and participated in the founding of Inference Corporation, which specializes in artificial intelligence. As a graduate student at California Technical Institute, Mr. Cole worked for IBM on mainframe and minicomputer software systems and was involved with the Goshawk Project. He is a member of numerous professional organizations and has authored several papers. Mr. Cole graduated magna cum laude from Harvard University, where he received bachelor's and master's degrees in Physics.

KORNEL TERPLAN, Ph.D.

Dr. Kornel Terplan is president of Performance Navigation Inc., an independent multinational communications consulting firm based in Hackensack, NJ and West Germany. During his career, he has been involved with the design of network operational control expert systems utilizing LISP and PROLOG machines, design of network management gateways for heterogeneous network architectures, and development and implementation of network- and mainframe-oriented reporting systems using SAS software. In his 15 successful years as a multinational consultant, Dr. Terplan provided training, consulting, and product development services to over 75 major national and international corporations. His book, *Communications Networks Management*, is considered the state-of-the-art abstract among international corporate users. Dr. Terplan received his doctoral degree at the University of Dresden. He completed advanced studies and research and lectured at Clemson University, Georgia Institute of Technology, Rensselaer Polytechnic Institute, and the University of California at Berkeley and Los Angeles.

For a complete listing of the topics covered in this service, see Index

Report	Report Number	No. of Pages	Pub'n. Date	Report	Report Number	No. of Pages	Pub'n. Date
INDEX (Tab NM01)							
Supplementary Index	NM01-050-101	2	8/89	Homegrown Network Management	NM30-200-401	4	6/89
Index	NM01-100-101	6	6/89				
USER'S GUIDE (Tab NM05)				INTEGRATED NETWORK MANAGEMENT (Tab NM40)			
User's Guide	NM05-100-101	6	6/89	Catching Up to the Future of Integrated Network Management	NM40-100-101	6	6/89
MANAGING NETWORKS (Tab NM10)				OSI-Based Network Management	NM40-200-101	16	6/89
Management by Preparedness	NM10-100-101	3	6/89	The OSI Network Management/Forum: A Critical Assessment	NM40-200-201	8	8/89
Network Management: A Manager's Perspective	NM10-100-301	8	6/89	Toward a Unified Theory of Managing Large Networks	NM40-200-401	5	6/89
Network Management: End-User Perspectives	NM10-100-401	3	6/89	Standardizing Network Management for TCP/IP Environments	NM40-300-101	4	6/89
Disciplines for Effective Network Management	NM10-200-301	9	6/89	Managing TCP/IP-based Networks: The HEMS Model	NM40-300-201	8	6/89
Evolving Market Opportunities in Network Management	NM10-300-101	6	6/89	AT&T Unified Network Management Architecture	NM40-313-101	14	6/89
NETWORK MANAGEMENT FUNCTIONS (Tab NM20)				Digital Equipment Corporation Enterprise Management Architecture (EMA)	NM40-325-101	9	6/89
Network Management Functions: Telecommunications Hardware	NM20-100-101	5	6/89	OpenView's Architectural Models	NM40-452-101	6	6/89
Fault Management				IBM SNA and NetView Network Management and Control Systems	NM40-491-101	8	6/89
Network Monitoring and Control	NM20-200-101	14	6/89		NM40-674-101	9	6/89
An Interactive Network Display System for Network Management Systems	NM20-200-201	10	6/89	NETWORK APPLICATIONS (Tab NM50)			
Configuration Management				SNA Networks			
Inventory and Configuration Management	NM20-300-101	4	6/89	IBM's Approach to Network Management	NM50-100-101	13	6/89
Change Control	NM20-300-201	7	6/89	Managing Change in SNA Networks	NM50-100-201	11	8/89
Performance Management				Network Management in APPN Networks	NM50-100-301	12	6/89
Performance Modeling: Analysis of Digital Communication Systems	NM20-400-101	11	6/89	T-Carrier and Digital-Based Networks			
Transport Management	NM20-400-201	4	6/89	T1 Network Management: A Strategic Perspective	NM50-200-101	4	6/89
Security Management				Choosing a T1 Network Monitoring System	NM50-200-201	2	6/89
Evolving Security Management Standards	NM20-500-101	7	6/89	T1 Multiplexers in the ISDN Environment	NM50-200-301	8	6/89
Disaster Recovery Planning in an Integrated Network Environment	NM20-500-201	13	6/89	Local Area Networks (LANs)			
Problem Management				Local Area Network Management Issues	NM50-300-101	8	6/89
Problem Management: Using the Help Desk	NM20-700-101	4	6/89	Token Bus/Ring LAN Management Concepts and Architecture	NM50-300-201	7	6/89
Designing a Practical Network Maintenance Strategy	NM20-700-201	7	6/89	Architectural Support of Network Management: An Alternate View	NM50-300-301	5	6/89
PLANNING & DESIGN (Tab NM30)				Managing Local Area Networks: Fault and Configuration Management	NM50-300-401	11	8/89
Future Scenarios for Network Management	NM30-100-101	4	6/89	Managing Local Area Networks: Accounting, Performance, and Security Management	NM50-300-501	8	6/89
Using Network Management Systems to Gain a Strategic Competitive Advantage	NM30-100-201	6	6/89				
Eight Critical Steps in Evaluating and Implementing Network Management Systems	NM30-200-101	14	6/89				
Designing Network Control Centers for Greater Productivity	NM30-200-201	6	6/89				

For a complete listing of the topics covered in this service, see Index

Report	Report Number	No. of Pages	Pub'n. Date	Report	Report Number	No. of Pages	Pub'n. Date
X.25 Packet Switched Networks				AT&T Network Management Strategy	NM60-046-101	5	6/89
Managing X.25 Packet Switched Networks	NM50-400-101	9	8/89	BellSouth Network Management Strategy	NM60-098-101	4	8/89
Circuit Switched Networks				IBM Network Management Strategy	NM60-491-101	8	6/89
Telecommunications Management Software	NM50-500-101	8	6/89	MCI Network Management Strategy	NM60-587-101	4	6/89
The Telemanagement Symphony	NM50-500-201	7	6/89	Network Equipment Technologies Network Management Strategy	NM60-660-101	4	6/89
Management of Centrex Systems	NM50-500-401	4	6/89	Northern Telecom Network Management Strategy	NM60-662-101	3	6/89
The Evolution of Automated Management Systems in Voice Networks	NM50-500-501	5	6/89	NYNEX Network Management Strategy	NM60-674-101	3	8/89
The Missing Link—Network Management	NM50-500-601	6	6/89	US SPRINT Network Management Strategy	NM60-895-101	3	6/89
Analog-Based Private Networks				SUPPLIERS & CONSULTANTS (Tab NM80)			
Modem/Multiplexer-Based Network Management	NM50-600-101	6	6/89	Directory of Suppliers	NM80-100-101	10	6/89
Network Management Systems for Data Communications	NM50-600-201	8	6/89	Directory of Consultants	NM80-200-101	35	6/89
NETWORK MANAGEMENT VENDOR STRATEGIES (Tab NM60)				GLOSSARY (Tab NM90)			
Strategies of Major Network Management Vendors	NM60-010-101	3	6/89	Glossary	NM90-100-101	79	6/89
User Evaluations of Vendor Offerings	NM60-010-201	4	6/89	NEWSLETTER (Tab NM99)			
				Communications Perspective		12	6/89

Supplementary Index

This Supplementary Index includes entries from the August issue which reflect additions or changes to the Index. To reference previously published reports, please consult the Index.

A

AT&T—
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201

B

Ballard, C.P.—
Managing Change in SNA Networks: NM50-100-201
Bell Operating Companies (BOCs)—
BellSouth Network Management Strategy: NM60-098-101
NYNEX Network Management Strategy: NM60-674-101
BellSouth—
BellSouth Network Management Strategy: NM60-098-101

C

change management—
Managing Change in SNA Networks: NM50-100-201
configuration management—
Managing Change in SNA Networks: NM50-100-201
Managing Local Area Networks: Fault and Configuration
Management: NM50-300-401
Corporation for Open Systems (COS)—
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201

D

data center management—
Managing Change in SNA Networks: NM50-100-201
Digital Equipment Corp.—
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201

F

Farfara, L.—
Managing Change in SNA Networks: NM50-100-201
fault management—
Managing Local Area Networks: Fault and Configuration
Management: NM50-300-401

H

Heldke, B.J.—
Managing Change in SNA Networks: NM50-100-201
Hewlett-Packard Co.—
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201

I

International Business Machines Corp.—
Managing Change in SNA Networks: NM50-100-201
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201

L

local area networks—
Managing Local Area Networks: Fault and Configuration
Management: NM50-300-401
local exchange carriers—
BellSouth Network Management Strategy: NM60-098-101
NYNEX Network Management Strategy: NM60-674-101

M

market projections—
BellSouth Network Management Strategy: NM60-098-101
NYNEX Network Management Strategy: NM60-674-101
Minoli, Dan—
Managing Local Area Networks: Fault and Configuration
Management: NM50-300-401

N

NetView—
Managing Change in SNA Networks: NM50-100-201
network applications—
Managing Change in SNA Networks: NM50-100-201
Managing Local Area Networks: Fault and Configuration
Management: NM50-300-401
Managing X.25 Packet Switched Networks: NM50-400-101
network management vendor strategies—
BellSouth Network Management Strategy: NM60-098-101
NYNEX Network Management Strategy: NM60-674-101
Northern Telecom Inc.—
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201
NYNEX—
NYNEX Network Management Strategy: NM60-674-101

O

OSI-based network management—
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201
OSI Network Management/Forum—
The OSI Network Management/Forum: A Critical Assessment:
NM40-200-201

Supplementary Index

P

packet switched networks—
 Managing X.25 Packet Switched Networks: NM50-400-101

S

Systems Network Architecture (SNA)—
 Managing Change in SNA Networks: NM50-100-201

V

value-added networks (VANs)—
 Managing X.25 Packet Switched Networks: NM50-400-101

vendors—

 BellSouth Network Management Strategy: NM60-098-101
 NYNEX Network Management Strategy: NM60-674-101

W

Wetterau, James B.—
 Managing X.25 Packet Switched Networks: NM50-400-101

X

X.25—
 Managing X.25 Packet Switched Networks: NM50-400-101 □

Index

A

- Abrahams, John B.—
 - Centrex-Based Network Management: NM50-500-401
- accounting management—
 - Eight Critical Steps in Evaluating and Implementing Network Management Systems: NM30-200-101
 - Network Management Functions: NM20-100-101
 - OSI-based Network Management: NM40-200-101
 - Standardizing Network Management for TCP/IP Environments: NM40-300-101
 - Token Bus/Ring LAN Management Concepts & Architectures: NM50-300-201
- ACCUMASTER integrator—
 - AT&T Unified Network Management Architecture: NM40-313-101
- Advanced-Peer-to-Peer Networking (APPN)—
 - Network Management in APPN Networks: NM50-100-301
- alerts—
 - IBM SNA and NetView: NM40-491-101
 - Network Management in APPN Networks: NM50-100-301
 - T1 Network Management: A Strategic Perspective: NM50-200-101
- analog-based private networks—
 - Modem/Multiplexer-based Network Management: NM50-600-101
 - Network Management Systems for Data Communications: NM50-600-201
- AT&T—
 - AT&T Network Management Strategy: NM60-046-101
 - AT&T Unified Network Management Architecture: NM40-313-101
 - Future Scenarios for Network Management: NM30-100-101
 - Homegrown Network Management: NM30-200-401
 - Strategies of Major Network Management Vendors: NM60-010-101
 - User Evaluations of Vendor Offerings: NM60-010-201

B

- Ball, Larry L.—
 - Network Management and Control Systems: NM40-674-101
- Barrett, Vince—
 - Future Scenarios for Network Management: NM30-100-101
- Bell Operating Companies (BOCs)—
 - Strategies of Major Network Management Vendors: NM60-010-101
- Ben-Artzi, Amatzia—
 - Standardizing Network Management for TCP/IP Environments: NM40-300-101
- Benhamou, Eric—
 - Standardizing Network Management for TCP/IP Environments: NM40-300-101
- Birdsong, Grady T.—
 - Choosing a T1 Network Monitoring System: NM50-200-201
- Brusil, Paul J.—
 - Toward a Unified Theory of Managing Large Networks: NM40-200-401

C

- call accounting—
 - Telecommunications Management Software: NM50-500-101
 - The Evolution of Automated Management Systems in Voice Networks: NM50-500-501
 - The Telemanagement Symphony: NM50-500-201

- Callon, R. W.—
 - Disciplines for Effective Network Management: NM10-200-301
- Centrex—
 - Centrex-Based Network Management: NM50-500-401
- change control—
 - Change Control: NM20-300-201
 - Inventory and Configuration Management: NM20-300-101
 - Requirements for a Network Management System Database: NM30-200-501
- change management—
 - Change Control: NM20-300-201
- Cincom Systems Inc.—
 - Strategies of Major Network Management Vendors: NM60-010-101
- circuit switched networks—
 - AT&T Unified Network Management Architecture: NM40-313-101
 - Centrex-Based Network Management: NM50-500-401
 - Eight Critical Steps in Evaluating and Implementing Network Management Systems: NM30-200-101
 - Network Management and Control Systems: NM40-674-101
 - Telecommunications Management Software: NM50-500-101
 - The Evolution of Automated Management Systems in Voice Networks: NM50-500-501
 - The Missing Link—Network Management: NM50-500-601
 - The Telemanagement Symphony: NM50-500-201
 - Transport Management: NM20-400-201
 - Using Network Management Systems to Gain a Strategic Competitive Advantage: NM30-100-201
- CLIST—
 - IBM SNA and NetView: NM40-491-101
 - IBM's Approach to Network Management: NM50-100-101
- Cole, Chris—
 - Management by Preparedness: NM10-100-101
- Common Management Information Protocol (CMIP)—
 - AT&T Unified Network Management Enterprise Architecture: NM40-313-101
 - Digital Equipment Corporation Enterprise Management Architecture (EMA): NM40-325-101
 - OSI-based Network Management: NM40-200-101
 - Standardizing Network Management for TCP/IP Environments: NM40-300-101
 - Toward a Unified Theory of Managing Large Networks: NM40-200-401
- Common Management Information Services (CMIS)—
 - AT&T Unified Network Management Architecture: NM40-313-101
 - Digital Equipment Corporation Enterprise Management Architecture (EMA): NM40-325-101
 - OSI-based Network Management: NM40-200-101
 - Toward a Unified Theory of Managing Large Networks: NM40-200-401
- configuration management—
 - Catching Up to the Future of Integrated Network Management: NM40-100-101
 - Change Control: NM20-300-201
 - Homegrown Network Management: NM30-200-401
 - Inventory and Configuration Management: NM20-300-101
 - Managing Local Area Networks: Configuration and Fault Management: NM50-300-401
 - Network Management: A Manager's Perspective: NM10-100-301
 - OSI-based Network Management: NM40-200-101
 - Standardizing Network Management for TCP/IP Environments: NM40-300-101
 - Token Bus/Ring LAN Management Concepts & Architectures: NM50-300-201
- Corker, K.—
 - Disciplines for Effective Network Management: NM10-200-301

Index

cost allocation—
Inventory and Configuration Management: NM20-300-101
CSUs/DSUs—
Architectural Support of Network Management: An Alternate
View: NM50-300-301

D

data center management—
Change Control: NM20-300-201
Inventory and Configuration Management: NM20-300-101
Problem Management: Using the Help Desk: NM20-700-101
Requirements for a Network Management System Database:
NM30-200-501
Digital Equipment Corp.—
Digital Equipment Corporation Enterprise Management
Architecture (EMA): NM40-325-101
Ethernet LAN Management: NMCC/VAX Ethernim, a Case Study:
NM50-300-601
Future Scenarios for Network Management: NM30-100-101
Homegrown Network Management: NM30-200-401
Strategies of Major Network Management Vendors: NM60-010-101
Digital Equipment Corp. Northern Telecom Inc.—
User Evaluations of Vendor Offerings: NM60-010-201
Digital Networking Architecture (DNA)—
Digital Equipment Corporation Enterprise Management
Architecture (EMA): NM40-325-101
Ethernet LAN Management: NMCC/VAX Ethernim, a Case Study:
NM50-300-601
director/entity model—
Digital Equipment Corporation Enterprise Management
Architecture (EMA): NM40-325-101
Ethernet LAN Management: NMCC/VAX Ethernim, a Case Study:
NM50-300-601
disaster recovery—
Disaster Recovery: NM20-500-201
Distributed Data Management (DDM)—
Network Management in APPN Networks: NM50-100-301
Dretler, John—
Eight Critical Steps in Evaluating and Implementing Network
Management Systems: NM30-200-101

E

equipment selection—
Choosing a T1 Network Monitoring System: NM50-200-201
Designing Network Control Centers for Greater Productivity:
NM30-200-201
Eight Critical Steps in Evaluating and Implementing Network
Management Systems: NM30-200-101
Using Network Management Systems to Gain a Strategic
Competitive Advantage: NM30-100-201
expert systems—
Architectural Support of Network Management: An Alternate
View: NM50-300-301
Catching Up to the Future of Integrated Network Management:
NM40-100-101
Network Management Systems for Data Communications:
NM50-600-201

F

fault management—
An Interactive Display System for Network Management Systems:
NM20-200-101
Choosing a T1 Network Monitoring System: NM50-200-201
Designing Network Control Centers for Greater Productivity:
NM30-200-201
Managing Local Area Networks: Accounting, Performance, and
Security Management: NM50-300-501
Managing Local Area Networks: Configuration and Fault
Management: NM50-300-401
Network Management Functions: NM20-100-101
Network Management in APPN Networks: NM50-100-301
Network Management: A Manager's Perspective: NM10-100-301
Network Monitoring and Control: NM20-200-101
OSI-based Network Management: NM40-200-101
Standardizing Network Management for TCP/IP Environments:
NM40-300-101

T1 Multiplexers in the ISDN Environment: NM50-200-301
T1 Network Management: A Strategic Perspective: NM50-200-101
Token Bus/Ring LAN Management Concepts & Architectures:
NM50-300-201

G

Gebremedhin, Elinor—
Network Management in APPN Networks: NM50-100-301
glossary—
Glossary: NM90-100-101
Goleniewski, Lillian—
The Telemanagement Symphony: NM50-500-201
government networks—
Architectural Support of Network Management: An Alternate
View: NM50-300-301
Disciplines for Effective Network Management: NM10-200-301
The Missing Link—Network Management: NM50-500-601
graphic interface—
An Interactive Display System for Network Management Systems:
NM20-200-101
Evolving Market Opportunities in Network Management:
NM10-300-101

H

help desk—
Designing a Practical Network Maintenance Strategy:
NM20-700-201
Problem Management: Using the Help Desk: NM20-700-101
Hewlett-Packard Co.—
Architectural Support of Network Management: An Alternate
View: NM50-300-301
Future Scenarios for Network Management: NM30-100-101
OpenView's Architectural Models: NM40-452-101
User Evaluations of Vendor Offerings: NM60-010-201
Higgerson, Clifford H.—
Evolving Market Opportunities in Network Management:
NM10-300-101
Horn, Robert J.—
Network Management Systems for Data Communications:
NM50-600-201

I

IBM NetView—
Catching Up to the Future of Integrated Network Management:
NM40-100-101
Future Scenarios for Network Management: NM30-100-101
Homegrown Network Management: NM30-200-401
IBM NetView and NetView PC—
Future Scenarios for Network Management: NM30-100-101
Homegrown Network Management: NM30-200-401
IBM's Approach to Network Management: NM50-100-101
Network Management in APPN Networks: NM50-100-301
IBM NetView/PC Version 1.1—
IBM's Approach to Network Management: NM50-100-101
integrated network management—
AT&T Unified Network Management Architecture: NM40-313-101
Catching Up to the Future of Integrated Network Management:
NM40-100-101
Digital Equipment Corporation Enterprise Management
Architecture (EMA): NM40-325-101
Disaster Recovery: NM20-500-201
Future Scenarios for Network Management: NM30-100-101
Homegrown Network Management: NM30-200-401
IBM SNA and NetView: NM40-491-101
Managing TCP/IP-based Networks: The HEMS Model:
NM40-300-201
Network Management and Control Systems: NM40-674-101
OpenView's Architectural Models: NM40-452-101
OSI-based Network Management: NM40-200-101
Standardizing Network Management for TCP/IP Environments:
NM40-300-101
Toward a Unified Theory of Managing Large Networks:
NM40-200-401
Interexchange carriers—
User Evaluations of Vendor Offerings: NM60-010-201

Index

International Business Machines Corp. (IBM)—

- Homegrown Network Management: NM30-200-401
- IBM Network Management Strategy: NM60-491-101
- IBM's Approach to Network Management: NM50-100-101
- Network Management in APPN Networks: NM50-100-301
- Strategies of Major Network Management Vendors: NM60-010-101
- User Evaluations of Vendor Offerings: NM60-010-201

internetworking—

- Architectural Support of Network Management: An Alternate View: NM50-300-301
- Managing TCP/IP-based Networks: The HEMS Model: NM40-300-201

inventory/configuration management—

- Change Control: NM20-300-201
- Inventory and Configuration Management: NM20-300-101
- Requirements for a Network Management System Database: NM30-200-501
- T1 Network Management: A Strategic Perspective: NM50-200-101

ISDN—

- Catching Up to the Future of Integrated Network Management: NM40-100-101
- T1 Multiplexers in the ISDN Environment: NM50-200-301
- User Evaluations of Vendor Offerings: NM60-010-201

J

Joseph, Celia—

- Network Management: A Manager's Perspective: NM10-100-301

Jurenko, Donald J.—

- The Missing Link—Network Management: NM50-500-601

K

Kanyuh, Denise—

- IBM's Approach to Network Management: NM50-100-101

Kappe, Laurie A.—

- Evolving Market Opportunities in Network Management: NM10-300-101

Klemba, Keith S.—

- OpenView's Architectural Models: NM40-452-101

Klessman, Horst—

- Local Area Network Management Issues: NM50-300-101

Kosak, Kerry—

- Designing Network Control Centers for Greater Productivity: NM30-200-201

Krentz, Dennis—

- Catching Up to the Future of Integrated Network Management: NM40-100-101

L

Llana, Andres Jr.—

- Disaster Recovery: NM20-500-201

local area networks (LANs)—

- Architectural Support of Network Management: An Alternate View: NM50-300-301
- Ethernet LAN Management: NMCC/VAX Ethernim, a Case Study: NM50-300-601
- Local Area Network Management Issues: NM50-300-101
- Managing Local Area Networks: Accounting, Performance, and Security Management: NM50-300-501
- Managing Local Area Networks: Configuration and Fault Management: NM50-300-401
- Standardizing Network Management for TCP/IP Environments: NM40-300-101
- Token Bus/ Ring LAN Management Concepts & Architectures: NM50-300-201

local exchange carriers—

- User Evaluations of Vendor Offerings: NM60-010-201

logical network management—

- IBM SNA and NetView: NM40-491-101
- Network Management: End-User Perspectives: NM10-100-401

M

mainstream network management—

- IBM SNA and NetView: NM40-491-101

managing networks—

- Disciplines for Effective Network Management: NM10-200-301
- Evolving Market Opportunities in Network Management: NM10-300-101
- Management by Preparedness: NM10-100-101
- Network Management: A Manager's Perspective: NM10-100-301
- Network Management: End-User Perspectives: NM10-100-401

market projections—

- Architectural Support of Network Management: An Alternate View: NM50-300-301
- AT&T Network Management Strategy: NM60-046-101
- Evolving Market Opportunities in Network Management: NM10-300-101
- IBM Network Management Strategy: NM60-491-101
- MCI Network Management Strategy: NM60-617-101
- Network Equipment Technologies (N.E.T.) Network Management Strategy: NM60-660-101
- Northern Telecom's Network Management Strategy: NM60-662-101
- Strategies of Major Network Management Vendors: NM60-010-101
- US SPRINT Network Management Strategy: NM60-895-101

McCann, John J.—

- OSI-based Network Management: NM40-200-101

McGarty, Terrence P.—

- Network Management and Control Systems: NM40-674-101

MCI Communications Corp.—

- MCI Network Management Strategy: NM60-617-101
- User Evaluations of Vendor Offerings: NM60-010-201

Meunier, J. M.—

- An Interactive Display System for Network Management Systems: NM20-200-101

Miller, Philip C.—

- Performance Modeling: Analysis of Digital Communication Systems: NM20-400-101

Minoli, Dan—

- Evolving Security Management Data Standards: NM20-500-101
- Managing Local Area Networks: Accounting, Performance, and Security Management: NM50-300-501
- Network Management Functions: NM20-100-101
- Network Monitoring and Control: NM20-200-101

modem-based network management—

- Architectural Support of Network Management: An Alternate View: NM50-300-301
- Modem/Multiplexer-based Network Management: NM50-600-101

modem management—

- Network Management Systems for Data Communications: NM50-600-201

monitoring and control—

- Choosing a T1 Network Monitoring System: NM50-200-201
- Homegrown Network Management: NM30-200-401
- Modem/Multiplexer-based Network Management: NM50-600-101
- Network Management and Control Systems: NM40-674-101
- Network Monitoring and Control: NM20-200-101
- The Evolution of Automated Management Systems in Voice Networks: NM50-500-501

Muller, Nathan J.—

- T1 Multiplexers in the ISDN Environment: NM50-200-301
- Telecommunications Management Software: NM50-500-101
- The Evolution of Automated Management Systems in Voice Networks: NM50-500-501
- Transport Management: NM20-400-201
- Using Network Management Systems to Gain a Strategic Competitive Advantage: NM30-100-201

multiplexer-based network management—

- Architectural Support of Network Management: An Alternate View: NM50-300-301
- Modem/Multiplexer-based Network Management: NM50-600-101
- Performance Modeling: Analysis of Digital Communication Systems: NM20-400-101

multiplexer management—

- Network Management Systems for Data Communications: NM50-600-201

Muralidhar, Kurudi K.—

- Network Management: A Manager's Perspective: NM10-100-301

Musselman, T.—

- T1 Network Management: A Strategic Perspective: NM50-200-101

Index

N

Neibaur, Dale—
Architectural Support of Network Management: An Alternate View: NM50-300-301

Net/Master (Cincom Systems)—
IBM SNA and NetView: NM40-491-101

NetView—
IBM SNA and NetView: NM40-491-101
User Evaluations of Vendor Offerings: NM60-010-201

NetView/PC—
IBM SNA and NetView: NM40-491-101

network applications—
Architectural Support of Network Management: An Alternate View: NM50-300-301
Centrex-Based Network Management: NM50-500-401
Choosing a T1 Network Monitoring System: NM50-200-201
Ethernet LAN Management: NMCC/VAX Ethernim, a Case Study: NM50-300-601
IBM's Approach to Network Management: NM50-100-101
Local Area Network Management Issues: NM50-300-101
Managing Local Area Networks: Accounting, Performance, and Security Management: NM50-300-501
Managing Local Area Networks: Configuration and Fault Management: NM50-300-401
Modem/Multiplexer-based Network Management: NM50-600-101
Network Management in APPN Networks: NM50-100-301
Network Management Systems for Data Communications: NM50-600-201
T1 Multiplexers in the ISDN Environment: NM50-200-301
T1 Network Management: A Strategic Perspective: NM50-200-101
Telecommunications Management Software: NM50-500-101
The Evolution of Automated Management Systems in Voice Networks: NM50-500-501
The Missing Link—Network Management: NM50-500-601
The Telemanagement Symphony: NM50-500-201
Token Bus/Ring LAN Management Concepts & Architectures: NM50-300-201

network control centers—
An Interactive Display System for Network Management Systems: NM20-200-101
Designing Network Control Centers for Greater Productivity: NM30-200-201
Disciplines for Effective Network Management: NM10-200-301
Network Management Systems for Data Communications: NM50-600-201

Network Equipment Technologies (N.E.T.)—
Network Equipment Technologies (N.E.T.) Network Management Strategy: NM60-660-101
User Evaluations of Vendor Offerings: NM60-010-201

network management architectures—
AT&T Unified Network Management Architecture: NM40-313-101
Digital Equipment Corporation Enterprise Management Architecture (EMA): NM40-325-101
IBM SNA and NetView: NM40-491-101

network management control—
Evolving Market Opportunities in Network Management: NM10-300-101

network management control and diagnostics (NMCD)—
Evolving Market Opportunities in Network Management: NM10-300-101

network management database—
Change Control: NM20-300-201
Inventory and Configuration Management: NM20-300-101
Problem Management: Using the Help Desk: NM20-700-101
Requirements for a Network Management System Database: NM30-200-501

network management functions—
An Interactive Display System for Network Management Systems: NM20-200-101
Change Control: NM20-300-201
Designing a Practical Network Maintenance Strategy: NM20-700-201
Disaster Recovery: NM20-500-201
Emerging Security Management Data Standards: NM20-500-101
Inventory and Configuration Management: NM20-300-101
Network Management Functions: NM20-100-101
Network Monitoring and Control: NM20-200-101
Performance Modeling: Analysis of Digital Communication Systems: NM20-400-101
Problem Management: Using the Help Desk: NM20-700-101

Transport Management: NM20-400-201

Network Management Protocol (NMP)—
AT&T Unified Network Management Architecture: NM40-313-101

network management vendor strategies—
AT&T Network Management Strategy: NM60-046-101
IBM Network Management Strategy: NM60-491-101
MCI Network Management Strategy: NM60-617-101
Network Equipment Technologies (N.E.T.) Network Management Strategy: NM60-660-101
Northern Telecom's Network Management Strategy: NM60-662-101
Strategies of Major Network Management Vendors: NM60-010-101
US SPRINT Network Management Strategy: NM60-895-101
User Evaluations of Vendor Offerings: NM60-010-201

network management vendors—
Ethernet LAN Management: NMCC/VAX Ethernim, a Case Study: NM50-300-601
IBM SNA and NetView: NM40-491-101

NMVT—
IBM SNA and NetView: NM40-491-101
IBM's Approach to Network Management: NM50-100-101
Network Management in APPN Networks: NM50-100-301

Nodine, M.—
Disciplines for Effective Network Management: NM10-200-301

Northern Telecom Inc.—
Future Scenarios for Network Management: NM30-100-101
Northern Telecom's Network Management Strategy: NM60-662-101
Strategies of Major Network Management Vendors: NM60-010-101

O

Ong, J.—
Disciplines for Effective Network Management: NM10-200-301

Open Systems Interconnect (OSI)—
Local Area Network Management Issues: NM50-300-101
Network Management and Control Systems: NM40-674-101
OSI-based Network Management: NM40-200-101

OSI-based network management—
AT&T Unified Network Management Architecture: NM40-313-101
Emerging Security Management Data Standards: NM20-500-101
Network Management Functions: NM20-100-101
Network Management: End-User Perspectives: NM10-100-401
OpenView's Architectural Models: NM40-452-101
OSI-based Network Management: NM40-200-101
Standardizing Network Management for TCP/IP Environments: NM40-300-101
Strategies of Major Network Management Vendors: NM60-010-101
Toward a Unified Theory of Managing Large Networks: NM40-200-401

OSI management framework—
Local Area Network Management Issues: NM50-300-101
OpenView's Architectural Models: NM40-452-101
OSI-based Network Management: NM40-200-101

OSI Management Information Base (MIB)—
OpenView's Architectural Models: NM40-452-101

OSI Network Management/Forum—
AT&T Unified Network Management Architecture: NM40-313-101
Digital Equipment Corporation Enterprise Management Architecture (EMA): NM40-325-101

OSI support—
AT&T Unified Network Management Architecture: NM40-313-101
Digital Equipment Corporation Enterprise Management Architecture (EMA): NM40-325-101
IBM SNA and NetView: NM40-491-101

P

packet switched networks—
Architectural Support of Network Management: An Alternate View: NM50-300-301
Disciplines for Effective Network Management: NM10-200-301

Papageorgiou, Chuck—
Homegrown Network Management: NM30-200-401

Partridge, Craig—
Managing TCP/IP-based Networks: The HEMS Model: NM40-300-201

PBX—
Eight Critical Steps in Evaluating and Implementing Network Management Systems: NM30-200-101

Index

- T1 Multiplexers in the ISDN Environment: NM50-200-301
 Telecommunications Management Software: NM50-500-101
 User Evaluations of Vendor Offerings: NM60-010-201
- PBXs—
 The Evolution of Automated Management Systems in Voice Networks: NM50-500-501
- performance management—
 Local Area Network Management Issues: NM50-300-101
 Managing Local Area Networks: Accounting, Performance, and Security Management: NM50-300-501
 Network Management Functions: NM20-100-101
 Network Management: A Manager's Perspective: NM10-100-301
 Network Monitoring and Control: NM20-200-101
 OSI-based Network Management: NM40-200-101
 Performance Modeling: Analysis of Digital Communication Systems: NM20-400-101
 Standardizing Network Management for TCP/IP Environments: NM40-300-101
 The Evolution of Automated Management Systems in Voice Networks: NM50-500-501
 Token Bus/Ring LAN Management Concepts & Architectures: NM50-300-201
 Transport Management: NM20-400-201
 Using Network Management Systems to Gain a Strategic Competitive Advantage: NM30-100-201
- performance modeling—
 Inventory and Configuration Management: NM20-300-101
- physical network management—
 AT&T Unified Network Management Architecture: NM40-313-101
 Homegrown Network Management: NM30-200-401
 Modem/Multiplexer-based Network Management: NM50-600-101
 Network Management: End-User Perspectives: NM10-100-401
- planning and design—
 Designing Network Control Centers for Greater Productivity: NM30-200-201
 Eight Critical Steps in Evaluating and Implementing Network Management Systems: NM30-200-101
 Future Scenarios for Network Management: NM30-100-101
 Homegrown Network Management: NM30-200-401
 Management by Preparedness: NM10-100-101
 Requirements for a Network Management System Database: NM30-200-501
 Using Network Management Systems to Gain a Strategic Competitive Advantage: NM30-100-201
- problem management—
 Designing a Practical Network Maintenance Strategy: NM20-700-201
 Management by Preparedness: NM10-100-101
 Problem Management: Using the Help Desk: NM20-700-101
 Requirements for a Network Management System Database: NM30-200-501
- Public Switched Telephone Networks (PSTN)—
 Centrex-Based Network Management: NM50-500-401
- Pyykkonen, Martin—
 Network Management: End-User Perspectives: NM10-100-401

R

- Rao, Anand V.—
 Designing a Practical Network Maintenance Strategy: NM20-700-201
- Remote Operation Services Element (ROSE)—
 AT&T Unified Network Management Architecture: NM40-313-101
 OSI-based Network Management: NM40-200-101
 Standardizing Network Management for TCP/IP Environments: NM40-300-101
- response time—
 Performance Modeling: Analysis of Digital Communication Systems: NM20-400-101
- REXX—
 IBM SNA and NetView: NM40-491-101
- Robertson, Jim—
 Standardizing Network Management for TCP/IP Environments: NM40-300-101

S

- Salazar, Andres C.—
 Network Management Systems for Data Communications: NM50-600-201

- Saydam, Tuncay—
 Token Bus/ Ring LAN Management Concepts & Architectures: NM50-300-201
- Scarfo, Philip J.—
 Network Management Systems for Data Communications: NM50-600-201
- security management—
 Disaster Recovery: NM20-500-201
 Emerging Security Management Data Standards: NM20-500-101
 Network Management Functions: NM20-100-101
 OSI-based Network Management: NM40-200-101
 Token Bus/Ring LAN Management Concepts & Architectures: NM50-300-201
- Sethi, Adarshpas, S.—
 Token Bus/Ring LAN Management Concepts & Architectures: NM50-300-201
- Shilling, Gary D.—
 Performance Modeling: Analysis of Digital Communication Systems: NM20-400-101
- sidestream network management—
 AT&T Unified Network Management Architecture: NM40-313-101
- Sligh, Robert L., Jr.—
 The Missing Link—Network Management: NM50-500-601
- SNA—
 IBM's Approach to Network Management: NM50-100-101
 SNA (System Network Architecture)—
 IBM SNA and NetView: NM40-491-101
 Network Management and Control Systems: NM40-674-101
- SNA networks—
 IBM's Approach to Network Management: NM50-100-101
 Network Management in APPN Networks: NM50-100-301
- source control—
 Change Control: NM20-300-201
- standards organizations—
 OSI-based Network Management: NM40-200-101
 Standardizing Network Management for TCP/IP Environments: NM40-300-101
- Stillman, M.—
 Disciplines for Effective Network Management: NM10-200-301
- Stokesberry, Daniel P.—
 Toward a Unified Theory of Managing Large Networks: NM40-200-401
- suppliers & consultants—
 Directory of Consultants: NM80-100-201
 Directory of Suppliers: NM80-100-101
- system implementation—
 Architectural Support of Network Management: An Alternate View: NM50-300-301
 Eight Critical Steps in Evaluating and Implementing Network Management Systems: NM30-200-101
- Systems Network Architecture (SNA)—
 Network Management in APPN Networks: NM50-100-301

T

- T-Carrier and Digital-Based Networks—
 Choosing a T1 Network Monitoring System: NM50-200-201
 T1 Multiplexers in the ISDN Environment: NM50-200-301
 T1 Network Management: A Strategic Perspective: NM50-200-101
- T1 multiplexer—
 User Evaluations of Vendor Offerings: NM60-010-201
- T1 network management—
 Choosing a T1 Network Monitoring System: NM50-200-201
 Modem/Multiplexer-based Network Management: NM50-600-101
 T1 Multiplexers in the ISDN Environment: NM50-200-301
 T1 Network Management: A Strategic Perspective: NM50-200-101
 Transport Management: NM20-400-201
- TCP/IP—
 Managing TCP/IP-based Networks: The HEMS Model: NM40-300-201
 Standardizing Network Management for TCP/IP Environments: NM40-300-101
- telecommunications management—
 T1 Multiplexers in the ISDN Environment: NM50-200-301
- telecommunications management software—
 Eight Critical Steps in Evaluating and Implementing Network Management Systems: NM30-200-101
 Telecommunications Management Software: NM50-500-101
 The Telemanagement Symphony: NM50-500-201
- telemanagement—
 The Telemanagement Symphony: NM50-500-201

Index

Token Bus—
 Token Bus/Ring LAN Management Concepts & Architectures:
 NM50-300-201
Token Ring—
 Token Bus/Ring LAN Management Concepts & Architectures:
 NM50-300-201
traffic analysis—
 The Evolution of Automated Management Systems in Voice
 Networks: NM50-500-501
transport management—
 Choosing a T1 Network Monitoring System: NM50-200-201
 Modem/Multiplexer-based Network Management: NM50-600-101
 Performance Modeling: Analysis of Digital Communication
 Systems: NM20-400-101
 T1 Multiplexers in the ISDN Environment: NM50-200-301
 The Evolution of Automated Management Systems in Voice
 Networks: NM50-500-501
 Transport Management: NM20-400-201
 Using Network Management Systems to Gain a Strategic
 Competitive Advantage: NM30-100-201
Trewitt, Glenn—
 Managing TCP/IP-based Networks: The HEMS Model:
 NM40-300-201
Tschammer, Voler—
 Local Area Network Management Issues: NM50-300-101

U

US SPRINT Communications Co.—
 US SPRINT Network Management Strategy: NM60-895-101
US West Network Systems Inc. (US WNSI)—
 Strategies of Major Network Management Vendors: NM60-010-101
use of this publication—
 User's Guide: NM05-100-101
user training—
 Designing a Practical Network Maintenance Strategy:
 NM20-700-201

V

Van der Meer, Roland A.—
 Evolving Market Opportunities in Network Management:
 NM10-300-101

Van Tijn, Judy—
 Problem Management: Using the Help Desk: NM20-700-101
vendors—
 Architectural Support of Network Management: An Alternate
 View: NM50-300-301
 AT&T Network Management Strategy: NM60-046-101
 Evolving Market Opportunities in Network Management:
 NM10-300-101
 IBM Network Management Strategy: NM60-491-101
 MCI Network Management Strategy: NM60-617-101
 Network Equipment Technologies (N.E.T.) Network Management
 Strategy: NM60-660-101
 Northern Telecom's Network Management Strategy:
 NM60-662-101
 Strategies of Major Network Management Vendors: NM60-010-101
 US SPRINT Network Management Strategy: NM60-895-101
virtual private networks—
 User Evaluations of Vendor Offerings: NM60-010-201

W

Wells, Andrea—
 The Telemanagement Symphony: NM50-500-201
Westcott, J.—
 Disciplines for Effective Network Management: NM10-200-301
wide area networks—
 Architectural Support of Network Management: An Alternate
 View: NM50-300-301
Wolpin, Stuart—
 Change Control: NM20-300-201
Woolston, Dayle S.—
 Architectural Support of Network Management: An Alternate
 View: NM50-300-301

X

X.500—
 Emerging Security Management Data Standards: NM20-500-101 □

User's Guide—How to Get the Most from Your Subscription

This report will help you to:

- Reap added value from *Datapro Network Management* by using it to its fullest.
 - Know the contents of each component of your subscription.
 - Take best advantage of the loose-leaf design features, structure, and organization.
-
-

We know that you are eager to explore *Datapro Network Management*. But before plunging in, take a few minutes now to familiarize yourself with your new information service through this handy User's Guide. It will quickly take you through the various components included in your subscription; point out the design features, structure, and organization of the loose-leaf binder; and show numerous ways that your subscription can serve you and your associates.

WHY DATAPRO NETWORK MANAGEMENT?

Simply put, network management is an activity too complex and crucial to be operated in a vacuum. Why not profit from the opportunities or problems that other network management professionals have already addressed or solved? Datapro's network management editors continuously review hundreds of challenging issues, technologies, and techniques and share the most successful solutions and ideas with our subscribers by presenting them in this service. *Datapro Network Management* can help you to manage existing operations successfully and also to spot emerging issues and trends, thus supporting your ability to anticipate opportunities, avoid pitfalls, and plan and control change.

Datapro Network Management carries the freshest insights and most viable approaches to the network management environment, both today and tomorrow. Having these reports at your fingertips is comparable

to having dozens of experienced network management specialists on call, with the added value of Datapro's direction to assure their focus on the pertinent issues facing network managers. Reports are presented in an easy-to-read, standardized format that makes it easy to find the specific information you need.

WHAT DOES YOUR SUBSCRIPTION INCLUDE?

Loose-Leaf Information Book

Like other Datapro services, the most visible component of your new *Datapro Network Management* subscription is its ever-changing compilation of loose-leaf reports. Designed for network management system planners and users, the reports are organized in logical progression—starting with a briefing on the basic issues, continuing with an in-depth study of network management functions, investigating integrated management concepts, focusing in on specific network management applications, and—finally—examining vendors' strategies.

Bimonthly Loose-Leaf Issues

Throughout your subscription term, *Datapro Network Management's* contents will steadily change and grow to help you with whatever problems or opportunities

User's Guide—How to Get the Most from Your Subscription

you encounter. We'll send you fresh new management reports bimonthly plus a monthly newsletter.

The Table of Contents and Index are continually updated to keep you current with the reports contained in your binder. You'll find more information on how to use these handy locator aids in the "How to Find Things Fast!" section below.

Monthly News Analysis

Maximize your investment in *Datapro Network Management* by circulating the monthly newsletter to those in your organization who need to know about communications issues. As you review each month's newsletter, you will likely find specific articles worth pointing out to others. By keeping your associates up to date on the ongoing value of the information contained in *Datapro Network Management*, you can multiply your subscription's value many times over.

HOW TO FIND THINGS FAST!

Datapro Network Management makes it easy for you to find information, whether you need a briefing on general concepts, an answer to a specific question, management tips to help you solve an impending problem, or "how to" guidelines on taking action. Become familiar with the major locator aids listed below:

The Table of Contents. The Table of Contents is located at the front of your service.

Updated with each bimonthly issue, the Table of Contents lists every report in page number order. Section headers, subheads, and report numbers help you to see the logical structure by which the reports are organized. The Table of Contents is useful for identifying clusters of reports that address a specific subject and provides a handy tool for quickly locating reports with which you are already familiar.

The Index. The Index is located in the Index section (Tab NM01). Every keyword associated with individual reports is listed in the Index, followed by an alphabetical listing by title of each report having that keyword. Report numbers are also provided for easy lookup.

Our industry frequently has several synonyms for the same concept, for example: *Network Control Centers (NCCs)* and *Network Operations Centers (NOCs)*. To avoid presenting an unnecessarily complex index, we've selected one preferred or most frequently used keyword for each concept and listed thereunder all related report titles. All synonyms for that keyword are

datapro NETWORK MANAGEMENT				Contents June 1989			
For a complete listing of the topics covered in this service, see Index							
Report	Report Number	No. of Pages	Pub'n. Date	Report	Report Number	No. of Pages	Pub'n. Date
INDEX (Tab NM01)				INTEGRATED NETWORK MANAGEMENT (Tab NM40)			
Index	NM01-100-101	20	6/89	Catching Up to the Future of Integrated Network Management	NM40-100-101	6	6/89
USER'S GUIDE (Tab NM05)				OSI-Based Network Management	NM40-200-101	16	6/89
User's Guide	NM05-100-101	7	6/89	Toward a Unified Theory for Managing Large Networks	NM40-200-401	5	6/89
MANAGING NETWORKS (Tab NM10)				Standardizing Network Management for TCP/IP Environments	NM40-300-101	4	6/89
Management by Preaddress	NM10-100-101	3	6/89	Managing TCP/IP-based Networks: The HEMS Model	NM40-300-201	8	6/89
Network Management: A Manager's Perspective	NM10-100-301	8	6/89	AT&T Unified Network Management Architecture	NM40-313-101	14	6/89
Network Management: End-User Perspectives	NM10-100-401	3	6/89	Digital Equipment Corporation Enterprise Management Architecture (EMA)	NM40-325-101	9	6/89
Disciplines for Effective Network Management	NM10-200-301	9	6/89	OpenView's Architectural Models	NM40-452-101	6	6/89
Evolving Market Opportunities in Network Management	NM10-300-101	6	6/89	IBM SNA and NetView Network Management and Control Systems	NM40-491-101 NM40-674-101	8 9	6/89 6/89
NETWORK MANAGEMENT FUNCTIONS (Tab NM20)				NETWORK APPLICATIONS (Tab NM50)			
Network Management Functions	NM20-100-101	5	6/89	SNA Networks			
Fault Management				IBM's Approach to Network Management	NM50-100-101	13	6/89
Network Monitoring and Control	NM20-200-101	8	6/89	Network Management in APPN Networks	NM50-100-301	12	6/89
An Interactive Display System for Network Management Systems	NM20-200-201	10	6/89	T-Carrier and Digital-Based Networks			
Configuration Management				T1 Network Management: A Strategic Perspective	NM50-200-101	4	6/89
Inventory and Configuration Management	NM20-300-101	4	6/89	Choosing a T1 Network Monitoring System	NM50-200-201	2	6/89
Change Control	NM20-300-201	8	6/89	T1 Multiplexers in the ISDN Environment	NM50-200-301	8	6/89
Performance Management				Local Area Networks (LANs)			
Performance Modeling: Analysis of Digital Communication Systems	NM20-400-101	12	6/89	Local Area Network Management Issues	NM50-300-101	8	6/89
Transport Management	NM20-400-201	4	6/89	Token Bus/Ring LAN Management: Concepts & Architectures	NM50-300-201	7	6/89
Security Management				Architectural Support of Network Management: An Alternate View	NM50-300-301	5	6/89
Evolving Security Management Standards	NM20-500-101	7	6/89	Managing Local Area Networks: Accounting, Performance, and Security Management	NM50-300-501	7	6/89
Disaster Recovery	NM20-500-201	13	6/89	Circuit Switched Networks			
Problem Management				Telecommunications Management Software	NM50-500-101	8	6/89
Problem Management: Using the Help Desk	NM20-700-101	4	6/89	The Telemanagement Symphony	NM50-500-201	7	6/89
Designing a Practical Network Maintenance Strategy	NM20-700-201	7	6/89	Centrex-Based Network Management	NM50-500-401	4	6/89
PLANNING & DESIGN (Tab NM30)				The Evolution of Automated Management Systems in Voice Networks	NM50-500-501	5	6/89
Future Scenarios for Network Management	NM30-100-101	4	6/89	The Missing Link—Network Management	NM50-500-601	7	6/89
Using Network Management Systems to Gain a Strategic Competitive Advantage	NM30-100-201	6	6/89				
Eight Critical Steps in Evaluating and Implementing Network Management Systems	NM30-200-101	14	6/89				
Designing Network Control Centers for Greater Productivity	NM30-200-201	6	6/89				
Homegrown Network Management	NM30-200-401	4	6/89				

COPYRIGHT © 1989 MCGRAW-HILL, INCORPORATED. REPRODUCTION PROHIBITED
DATAPRO RESEARCH, DELRAN NJ 08075 USA

Figure 1. The Table of Contents is updated with each issue and lists every report in the order in which it appears in the volume. Section headers, subheads, and report numbers help you to see the logical structure within which the reports are organized. The Table of Contents is useful for identifying clusters of reports that fit a specific management need and is also a handy tool for quickly locating reports with which you are already familiar.

cross-referenced to lead you to the list of relevant titles, no matter what term you start with. Acronyms are also cross-referenced, and related subjects are grouped under one specific keyword. Authors' names, for reports written by industry authors, are also listed alphabetically in the Index.

The Index is updated two times per year. Between updates, a Supplementary Index references new keywords and titles introduced in the intervening issues.

The Index is your most powerful locator aid. By referring to it frequently, you will greatly increase your ability to find specific subjects of interest.

While you may enjoy occasional browsing, *Datapro Network Management* is not intended for leisurely reading. It is a workhorse designed to produce results. If you exercise this information service frequently, using the locator aids we've described, you can expect to

User's Guide—How to Get the Most from Your Subscription

reap hundreds of practical suggestions, cost-saving tips, and sound advice worth many times the cost of your subscription.

VENDOR STRATEGIES PROVIDED BY NBI/DATAPRO

The reports behind Tab NM60, "Network Management Vendor Strategies," are provided by Northern Business Information/Datapro (NBI/Datapro). NBI is Datapro's telecommunications market research arm, providing research and industry analysis on key markets, aftermarkets, services, and companies.

NBI/Datapro also offers a complete family of products for strategic planners in the telecommunications industry, both vendors and end users. Key industry segments are analyzed and delivered in a combination of reports, diskettes, newsletters, seminars, and limited access to NBI's database.

OTHER DATAPRO SOURCES OF INFORMATION ABOUT TELECOMMUNICATIONS

Related Loose-Leaf Services

Today's communications professionals are no longer confined to the narrow role of liaison between user organizations and the local telco representatives. They are now becoming accountable for forming long-range telecommunications strategies, recommending and procuring equipment and transmission facilities from numerous vendors in a highly volatile marketplace, and coordinating or integrating voice systems and networks with their organizations' data communications facilities.

No single information service can address all these topics, but Datapro has an answer! To complement *Datapro Network Management*, we offer nine other loose-leaf information services in our Communications Series. Together they form a complete library of information on both voice and data communications.

The other members of the set are:

- *Datapro Management of Telecommunications*. This management-oriented information service provides a comprehensive set of guidelines for researching, planning, designing, and managing voice communications networks.
- *Datapro Reports on Telecommunications*. This product-oriented companion to *Datapro Manage-*

ment of Telecommunications contains hundreds of authoritative reports, comparisons, and evaluations of telephone systems hardware and software, related voice equipment, and both basic and value-added transmission facilities. It is designed for use during the equipment selection stage of the telecommunications management process.

- *Datapro Management of Data Communications*. This management-oriented information service, formerly titled *Datapro Communications Solutions*, provides a comprehensive set of guidelines for planning, designing, and managing data communications networks. This service is similar in charter and in format to *Management of Telecommunications* but is targeted towards data communications managers.
- *Datapro Reports on Data Communications*. This information service addresses the data communications marketplace and provides in-depth descriptions; detailed comparison columns; and insightful commentary on commercially available data communications devices, systems, and services.

NM30-100-201
Planning & Design

Using Network Management Systems to Gain a Strategic Competitive Advantage

This report will help you to:

- Learn how major corporations use network management systems to increase and sustain their competitive advantage.
- Use the capabilities of your existing network management equipment to maximize your network's strategic value.

Traditionally, users viewed corporate networks as mere utilities that supported the organizational infrastructure by facilitating communications among far-flung locations. The control inherent in private facilities translated into substantial cost savings over the long term—thus justifying the expense of running one's own network.

In the early 1980s, industry experts promoted private networks as the means of making geographically dispersed companies more manageable—achieving cohesiveness among diverse operating units and helping top-heavy organizations trim the burgeoning ranks of middle management. Ideally, a private network would promote better and more timely decision making, translating into a more profitable business. The ideas had some merit, but for a long time the appeal of private networks remained focused on anticipated cost savings and on increasing the flexibility in allocating communications resources.

Today, a ground swell of opinion says that the quality of a company's network holds the key to serving customers better, increasing market share, and pursuing new business opportunities successfully. In the process, an enterprise can secure strategic competitive advantages or marketplace rivals that have not yet awakened to such possibilities.

This report was developed exclusively for Datapro by Nathan J. Muller, a former consultant. Mr. Muller has 15 years' experience in the computer and telecommunications industries. He has written extensively on all aspects of computers and communications and is the author of "Minimum Risk Strategy for Assuring Communications Equipment and Services" (Artech House, 1987).

Many accept this vision of the corporate network as a competitive weapon as self-evident. Without further explanation, however, one might be led to believe that the particular arrangement of lines or the type of equipment deployed among the various nodes determines the network's strategic value.

Although certainly essential, network architecture and components take a backseat to the network management system (NMS). The right NMS unifies diverse computer and communications resources and transforms them into strategic assets that improve a company's competitive position and long-term survivability. Not surprisingly, then, many organizations place a premium value on network downtime:

- A Wall Street brokerage house can lose as much as \$60,000 per minute when buy/sell instructions from customers are disrupted.
- A state lottery can lose millions of dollars per hour if it cannot process ticket sales when hotly jackpots are at stake.
- An insurance company can lose its *Fortune* 500 accounts if it cannot live up to specified levels of network uptime to process its clients' claims.

As more and more companies are discovering, the capability to maintain a strategic competitive advantage rests upon the quality of their networks. The "network" includes the high-capacity backbone, feeder links, and drops typically associated with wide area networks (WANs) as well as the host servers.

JUNE 1989 COPYRIGHT © 1989 MCGRAW-HILL, INCORPORATED. REPRODUCTION PROHIBITED
DATAPRO RESEARCH, DELRAN NJ 08075, USA

Figure 2. The first page of each report lists bulleted objectives that capsulize the report's contents. These results-oriented statements focus on what you should expect to obtain from reading the report.

User's Guide—How to Get the Most from Your Subscription

- *Datapro Reports on Communications Alternatives.* This service covers transmission techniques, networking systems, and products associated with T-carrier, fiber optic, microwave, satellite, and infrared systems.
- *Datapro Reports on PC Communications.* This product is a comprehensive, three-volume information service providing detailed analysis of PC-to-host links, asynchronous communications packages, modems, and local area networks based on hands-on tests. In addition, Technology Reports address innovative advances, and Survey Reports are provided to detail the entire microcomputer communications marketplace.
- *Datapro Reports on Communications Software.* Combines general industry and system overviews with product-specific reports in a fully indexed, loose-leaf format. This service also includes product reports on network management packages, such as IBM's NetView and Cincom's Net/Master.
- *Datapro Networking Services.* Analyzes telecommunications common carrier rate changes and offerings.
- *Datapro Directory of On-line Services*—features profiles of more than 2,000 publicly accessible databases and full feature reports on the major information retrieval and remote computing services.

In addition, we publish a set of services addressed specifically to international audiences. These include *Datapro Reports on International Telecommunications*, *Datapro Reports on Data Communications International*, and *Datapro Reports on International Communications Equipment*.

To make any of these information services a part of your reference library, you can first get a trial review copy by calling Datapro Customer Service or your Datapro account representative toll free at (800) DATAPRO (800/328-2776). Examine the review copy in your office for 30 days and agree to return the books and pay a small review fee if you decide not to subscribe.

The table below clearly shows the relationships between these services:

	<u>Orientation</u>	<u>Technology</u>
Datapro Management of Telecommunications	Management	Voice
Datapro Reports on Telecommunications	Products	Voice
Datapro Management of Data Communications	Management	Data
Datapro Reports on Data Communications	Products	Data
Datapro Reports on Communications Alternatives	Product/Management	Transmission
Datapro Network Management	Management	Network Management
Datapro Reports on PC Communications	Products	Data
Datapro Reports on Communications Software	Products	Data
Datapro Networking Services	Tariffs	Services
Datapro Directory of On-line Services	Products	Services

Other Loose-Leaf Services

Datapro Information Services include comprehensive loose-leaf subscription services, covering microcomputers, information processing systems, data and telecommunications, office automation, information security, and manufacturing automation. A list of Datapro products appears on the following page:

User's Guide—How to Get the Most from Your Subscription

INFORMATION SYSTEMS

Datapro 70
Datapro 70 International
Reports on Minicomputers
Reports on Minicomputers International
Reports on Information Security
Reports on Software
Reports on Software International
Reports on UNIX Systems & Software
Directory of Software
Management of Applications Software
Management of EDP Systems

MICROCOMPUTERS

Reports on Microcomputers
Reports on Microcomputers International
Management of Microcomputer Systems
Directory of Microcomputer Software
Directory of Microcomputer Hardware

OFFICE AUTOMATION

Reports on Office Automation
Reports on Office Automation International
Management of Office Automation
Reports on Word Processing
Reports on Word Processing International
Office Products Evaluation Service
Reports on Electronic Publishing Systems

INDUSTRY AUTOMATION

Reports on Banking Automation
Reports on Retail Automation
Manufacturing Automation Series
Management and Planning
Manufacturing Information Systems
CAD/CAM/CAE Systems
Factory Automation
News and Perspectives
Reports on Marketing Information Systems

Professional Development Seminars

Through Datapro, your personal development can extend well beyond the use of loose-leaf information services. We currently offer 19 different seminars in voice and data communications, networking, and connectivity. Datapro subscribers are treated as preferred customers who receive priority enrollment. Please call Datapro Customer Service for the latest course offerings.

We also offer a broad range of on-site courses that can save your company time and money. Courses cover topics such as communications, information systems, personal computing, software, and management. Call for free detailed outlines today.

Computer-Based Training (CBT)

Datapro offers state-of-the-art courseware packages providing interactive, animated, and graphical lessons that you learn at home or work on an IBM or compatible PC. Ten CBT disks feature training in various communications disciplines, including data communications, telecommunications, LANs, PC networking, X.25, and T1 transmission. A demonstration disk, available for each course, provides an overview with actual course excerpts.

Special Publications

Datapro publishes dozens of feature reports, soft-bound books, and loose-leaf books (*not* updated under subscription) on many topics, including telecommunications. Our Customer Service staff will gladly send you a copy of the latest catalog.

We also provide third-party, objective research data that can be tailored to suit an organization's needs. This highly customized information—available from the entire contents of the Datapro and Future Computing databases—includes the following products and services: Electronic Database Services; Custom Publishing and Market Research; Competitive Handbooks; Future Computing Incorporated—micro/PC hardware and software testing, and market/distribution channel analysis; On-Site Educational Services; and Reprints and Feature Reports.

CUSTOMER SERVICE— DON'T OVERLOOK IT!

While we are considered most frequently a publisher, Datapro Research is, first and foremost, a manager of information for our subscribers. Accordingly, we dedicate ourselves to customer service. Many of the ways in which our Customer Service staff can serve you have already been outlined. Additional services include:

- *Refurbishing.* Missing pages, lost issues, broken binders, and damaged books are readily replaced.
- *Reprints.* Our subscribers frequently request reprints of selected reports to share with their associates and customers. Datapro can quickly fill reprint orders for

User's Guide—How to Get the Most from Your Subscription

several copies or several thousand at reasonable prices. Note that Datapro's copyrights prohibit any reproduction, duplication, or storage (on any medium, including electronic) of any portion of our information services without our written permission.

- *Subscription-related services.* Our Customer Service staff will gladly assist you in initiating trial reviews, reinstating expired subscriptions, making subscriber name and/or address changes, consolidating multiple subscriptions into a single billing, and other subscription fulfillment services.
- *Archival reports and back issues.* Datapro maintains an archive of reports published since 1970. Subscribers interested in tracking product histories, market evolutions, and other issues can use this storehouse of information. Copies of old reports and/or assistance in searching Datapro's archives can be obtained through Customer Service. Costs and availability of such reports vary depending on the specific nature of your request.

To make room for new material, we instruct you from time to time in the Filing Instructions to replace or delete older reports from your binders. You may wish to set up your own informal system for archiving reports that have been removed. Datapro can provide you with additional binders for a nominal charge.

- *Special handling.* Customer Service has a variety of mechanisms to expedite delivery of Datapro information to meet special requirements, including courier services, facsimile, and special packaging. Inquire for details.

NOW IT'S UP TO YOU

The basic goal of this—or any other—conscientious information service is to synthesize all the data about a selected topic into a comprehensive, orderly structure that is timely, accurate, and easy to use. *Datapro Network Management*, far more than a simple reference book, is an information *service* constantly refreshed to provide ongoing value to communications professionals.

Our value to you evolves from two unique Datapro capabilities. First, as the undisputed leader in the business of supplying loose-leaf information services to the data processing, office automation, and communications industries, we are able to monitor these industries as an insider. We are equipped to track literally hundreds of new technologies and industry trends, analyze them with the insight gained through experience, and present them to our readers on an exceptionally timely basis. Second, since no environment, especially the communications arena, remains static, we have created *Datapro Network Management* as a loose-leaf service that is updated continually with fresh new ideas and analysis. Each component and feature of the service is designed to maximize the benefits you receive from your subscription.

Now it is up to you. Put the service to work and give yourself the opportunity to find out how valuable it can be! □

Management by Preparedness

This report will help you to:

- Anticipate increases in trouble calls as your network grows in size and complexity.
 - Assess the impact that a one-terminal-per-employee operation will have on your network management strategy.
 - Prepare a realistic plan for managing future network expansion and change.
-
-

Network managers cannot escape the existence of failures in their networks. Although an aggressive preventive maintenance program can minimize failure, it cannot prevent it. Unmanaged failure can turn into disaster; therefore, network managers have to anticipate failure.

How many trouble calls should you be receiving? For a small data processing network under 1,000 devices, 5 percent of the end users will call in complaints on any given day. In larger networks, that percentage goes up quite a bit and may reach almost 40 percent on absolutely huge networks. Clearly, complaints need to be reduced if networks are to grow.

It is important to realize that hardware failures are not the primary cause of most trouble calls, particularly in large networks. The potential sources of trouble are numerous. They may include problems with applications and systems software as well as problems caused by uncontrolled network changes (terminal moves, report revisions, source program modifications, etc.). Table 1 lists some potential sources for

trouble. Note the different job descriptions of staff members who typically wind up answering these trouble calls.

Furthermore, as the networks increase in size, they increase in complexity. While the potential for trouble calls increases quadratically as the network grows, network managers can avoid disaster by preparing to handle complexity and growth. The following section portrays a realistic picture of the amount of investment and attention demanded by a well-run network.

A WELL-RUN NETWORK

As an example of a well-run network, consider the public phone network. An average-size central office (CO) has about 50,000 incoming lines. That means it serves about 50,000 customers or at least 50,000 phones. Typically, central offices receive a trouble call from about 1/10 of one percent of those phones each day or approximately 50 calls a day. Contrast that failure rate (1/10 of a percent) to 5 percent in a small data network. Of course, the phone network is much simpler than the data network; the phone network has a simple handset in it and uses technology that has been around for a long, long time. It is extremely simple compared to data processing terminals and programs, and all the associated devices and processes.

This report was written exclusively for Datapro by Chris Cole. Mr. Cole has over 12 years' experience in the data communications industry, and particularly in the management of large networks. Mr. Cole is currently the president and cofounder of Peregrine Systems, Inc. and an advisor to *Datapro Network Management*.

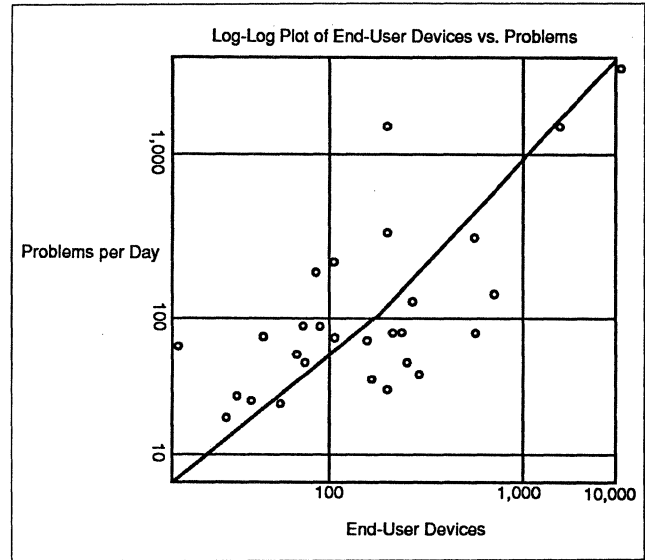
Management by Preparedness

How much is spent on preventive maintenance and error correction at a telephone CO? The typical staff of a CO is four people, two of whom are occupied full-time performing preventive maintenance. A third person is dedicated to correcting problems, and the fourth person is responsible for installations and changes to the system.

The telephone system's engineering has a built-in two for one redundancy. In other words, in the event of a component failure, there is a backup component that will activate. Despite that level of redundancy, a typical CO experiences at least one major alarm all the time. That means that there is at least one failed component with only one level of redundancy left—failure of the backup component would, in fact, cause a disruption of service. One maintenance worker is fixing failed components—day and night.

Phone equipment is very heavily overengineered. The electronic components are not run at anywhere near their tolerance. This is deliberate because the phone companies actively seek to minimize the failure rate in a switch. The planned failure rate in a switch is one hour of failure in every 25 years of operation. Even though the phone companies have massively over-engineered the switch, there is still one major alarm at any time.

How much does that overengineering cost? At least 40 percent of the switch's cost comes from redundancy and reliability features. These features, such as



Source: Peregrine Systems.

Figure 1. Network complexity increases the potential for trouble calls. Measuring complexity is difficult. One reasonable measure of complexity is the number of connections between components in the network. These connections may be physical (circuits), logical (call/dataflow), or procedural (approval groups/escalation paths). The number of connections increases as devices are added to the network. This figure depicts a log-log plot of problems per day vs the number of end-user devices.

self-checking and automatic and on-demand diagnostics, are all designed into the switch specifically to deal with failure.

In general, it is realistic to say that 70 percent of a phone bill is going towards the management of that network. This includes preparing for, preventing, and dealing with failure of that network. Most phone users think they are paying for electricity or copper wiring or the lease for the right-of-way phone lines when, in fact, that is a small fraction of the actual cost.

In the world of phone networks, which are very stable, network management accounts for about 70 percent of the network's cost. Network managers in the world of data processing do not recognize the cost of effective network management. To maintain a large, stable network with 50,000 terminals, a company must spend substantial amounts of money to keep the network running. Based on research, a 100,000-terminal network will generate 40,000 problems a day. (See Figure 1.) There are not many networks of that size, and those that are probably are not getting 40,000 trouble calls per day. The reason they receive fewer calls is not because there are fewer problems—it is because users have given up calling.

Problem	Person Receiving the Problem Call
On-line ABENDS	Operators
Batch ABENDS	Applications Programming
Lost reports	I/O or Production Control
VTAM vary ACTIVE	Network Control
TSO cancel user ID	Operators
Bugs (system)	Systems Programming
Bugs (application)	Applications Programming
B37 ABENDS	DASD Management
Data set lockouts	Operators
Terminal/line problems	Network Control
Forgotten password	Security
Tape problems	Tape Librarian
Scheduling problems	Operators or Production Control
PC problems	Infocenter
End-user help	Customer Support

Source: Peregrine Systems.

Table 1. Some typical categories of trouble calls. MIS managers in medium-to-large networks are rarely aware of the actual number of trouble calls they receive. This lack of awareness is because the calls are not logged and are not received by a central Help Desk.

Management by Preparedness

The consequences of network instability are devastating. The following two cases are real. While these cases may seem extreme, they are all too similar to the experiences of many users in large networks.

Case 1. An analyst for a large brokerage can remotely log on to the corporate mainframe to obtain information necessary to conduct business. After trying repeatedly, the analyst discovered he could not rely on the network. The analyst now uses a PC connected to a LAN; he would rather have a reliable PC than an unreliable mainframe terminal. Of course, the analyst has to live without the valuable corporate database. This user's experience is not unusual. It is safe to conclude that corporate America has spent a significant amount of money on PCs to avoid network management problems.

Case 2. A national corporation has a large network supporting hundreds of users. Every day at approximately 10:30 a.m., the network goes down until the technical support staff can restore service—a process which usually takes two or three hours. The root of the problem has yet to be solved. Over 300 programmers use the system for applications development, and they are literally out of work for two hours each day for lack of effective network management.

A LOOK AT THE FUTURE

There are about 10,000 networks in the world with 500 or more devices. The average size of these networks is about 1,000 devices. About 5,000 of those networks are in the United States. Those 5,000 networks are owned by the 5,000 largest companies in the country which, together, employ 50,000,000 people (or an average of 10,000 people per company).

U.S. corporations are moving towards a single terminal-per-employee operation. Over the next 5 to 10 years, the average network will grow from 1,000 devices to 10,000 devices. For example, Federal Express and UPS have installed terminals in each truck. Examples in other industries are not uncommon: there are now terminals on the shop floor of many manufacturing firms. Soon, there will be a terminal wherever your employee is—whether your business is a retail outlet or a manufacturing facility or a clerical group—every firm is evolving to a one-terminal-per-employee operation. This will drive the average network size from 1,000 terminals to 10,000 terminals. That will drive the average complaints per day from 50 calls to about 880 calls. Consequently, the cost of those calls will jump from under \$1 million to over \$10 million. Thus, while network managers today may be able to control their networks in a rather ad hoc fashion, that type of network management will not work in the future.

This is the lesson the phone company learned years ago, and that's when it discovered that it must spend over half of its money just on network management—and the technology that they are trying to manage is much simpler than data processing technology.

The point of these examples is that managers cannot just throw the network in and expect it to work. Companies must spend a substantial fraction of their budget—even the majority of it—on network management. This is the meaning of “management by preparedness,” and the industry is not ready for it. Managers must realize that for every \$10 million spent on a network, in the future \$7 million may be spent on what appears at first to be unproductive uses—network management, i.e., maintenance and prevention. Large networks simply will not run without that kind of expenditure. Management must get serious about preparing for failures. □

Network Management: A Manager's Perspective

This report will help you to:

- Understand the basic elements and functions of network management.
 - See how these functions might be used in a network management implementation.
 - Examine the current status and future outlook for network management standards and products.
 - Identify issues that managers should consider when planning to implement network management.
-

INTRODUCTION

The purpose of this report is to give an overview of network management from a technical manager's perspective. That is, we consider issues pertinent to a manager who is considering how to implement network management. These issues include such topics as how network management fits into an organizational environment, what products are available now, and where network management is going in the future. Where possible, we explain pertinent concepts in nontechnical terms.

As used in this report, network management refers to the management of data communications resources that use standard interfaces; i.e., open systems. Vendor proprietary network management schemes are outside our consideration. Although network management functions apply to all networks, some functions should be tailored to take the special needs of the applications running on the network into account. Throughout this report, our discussion is oriented to

the needs of local area networks (LANs), and we use the Manufacturing Automation Protocol (MAP) and Technical and Office Protocol (TOP) version 3.0 specifications^{6,7} of network management as a basis.

What is network management and why is it important? Network management can be defined as coordinating, monitoring, and controlling the distributed resources throughout a network. Network management is important because it can provide a wide range of information on a network, as well as a powerful set of tools for managing the network. For example, the types of information that network management can provide include network configuration and status, performance and trends, and current and pictorial operations data. Using network management tools, a network administrator can modify the network configuration, change its status, adjust parameters to tune performance, and analyze the location of and best solution for faults. As organizations rely more heavily on networks and networks become increasingly complex, management information and tools become critical to the organization's operation.

The report begins with background information on network management, including the elements needed to put it into place, and the basic functions it provides. Next, we give an example of network manager

This Datapro report is based on "Network Management: A Manager's Perspective," by Celia Joseph and Kurudi H. Muralidhar, Industrial Technology Institute, from *ENTERPRISE Conference Proceedings*. © 1988, Society of Manufacturing Engineers. Reprinted by permission.

Network Management: A Manager's Perspective

application; that is, given the set of basic functions, this is one example of how they could be used. We then examine the status of network management today: what products can do and what capabilities should be forthcoming. This is followed by our perception of where network management is heading in the long-range. We conclude with a short list of issues that a manager should consider before implementing a network management system.

BACKGROUND

Network Management Definition

A basic definition of network management is the controlling and supervising of the resources which allow communications to take place over the LAN.⁴ The term "resources" typically refers to network components (workstations, file servers, etc.), although it may also refer to the resources within a component, such as protocol layers or any configurable parameters that the device may have.

More specifically, network management includes ensuring the correct operation of the LAN, monitoring the use of network resources, maintaining network components in good working order, planning for changes to the network, and producing a variety of information on network operations—such as periodic or ad hoc reports.⁶

Network Management Environment

Network management is frequently viewed as only a technical problem. However, the ultimate responsibility for management resides with people, not machines. For example, a number of people within an organization will be involved with network management, including the users of the network, who may need access to current LAN status information; the managers throughout the organization, who may be concerned about how the LAN's performance will impact the performance of their portions of the organization; and the actual network administrator, who is in charge of the day-to-day operation of the LAN. Thus, the integrated management environment within which network management resides is a combination of human, social, and organizational resources, in addition to technology.

Figure 1 shows the main components of this integrated management environment. These include the users of the network, the network and systems resources, the manufacturing enterprise or company management, and network management.

- **Users of the network** are those interested in the operation and utilization of the network.
- **Network and system resources** consist of end systems (workstations, controllers, file servers, etc.), relay systems (bridges, routers, and gateways), and network components—in other words, the collection of objects that require managing.
- **Manufacturing enterprise management** directly impacts the network management functions by setting policies for the organization, such as whether Open Systems Interconnection (OSI) concepts will be used for communications, and what type of management structure (centralized, distributed) will be used.
- **Network management** consists of a combination of human, software, and hardware elements. The human element consists of the network administrators who make decisions on network management. The software and hardware represent the automated network manager tool that provides capabilities for the manufacturing enterprise network.

Network Management Elements

Given the overall environment for network management, we now take a closer look at how network management could be added to a company's communications environment. Figure 2 shows an example of a LAN with network management elements added. Each device has a LAN component that permits it to attach to the LAN. In this example, all of the LAN

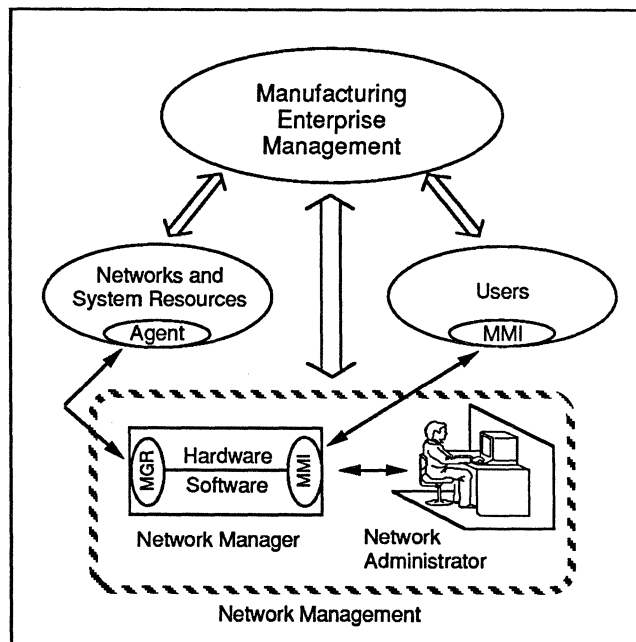


Figure 1. Integrated management environment.

Network Management: A Manager's Perspective

components use the standard, layered protocols specified by MAP/TOP in their interfaces. Some devices may share a LAN component. As defined by the standards groups, network management "fits" into and around the layered protocols in the LAN interfaces. However, more is needed to put network management into place than an element in a LAN component's interface. The full set of elements needed to put network management into place is as follows:

- **Network Administrator:** The person or persons who use the Network Manager to perform network management functions.
- **Network Manager Application:** This is an automated tool with a special man-machine interface (MMI) that the network administrator (the person) uses to monitor and control LAN activities. The LAN may have more than one network manager application.
- **Agent-SMAPs and Manager-SMAPs:** The agent-system management application process (Agent-SMAP) is a program that resides in each LAN component and manages the resources within the LAN component, as well as communicates with the Manager-SMAP. The Manager-SMAP is an analogous program that resides in the network manager application's LAN component.

- **Network Management Protocol:** The network management protocol is a set of rules that define how a Manager-SMAP communicates with Agent-SMAPs. This protocol is sometimes called the Manager-Agent protocol.
- **Management Information Base:** The management information base (MIB) consists of the information on resources that is maintained in each network device; that is, each device maintains information on its own resources. In the figure, the MIB is included in the "NM" element for each device. In addition, the network manager will maintain an information base for the domain for which it is responsible; this is included in the box marked "NM functions."
- **Management Domain:** A management domain is the set of all Agent-SMAPs which report to the same Manager-SMAP, or in other words, the set of devices that a network manager application will manage. If the LAN has a single network manager application, then all devices will be in that manager's domain. However, a LAN may have multiple manager applications. In this case, a domain must be defined for each manager and an agent may be in more than one manager's domain.⁶

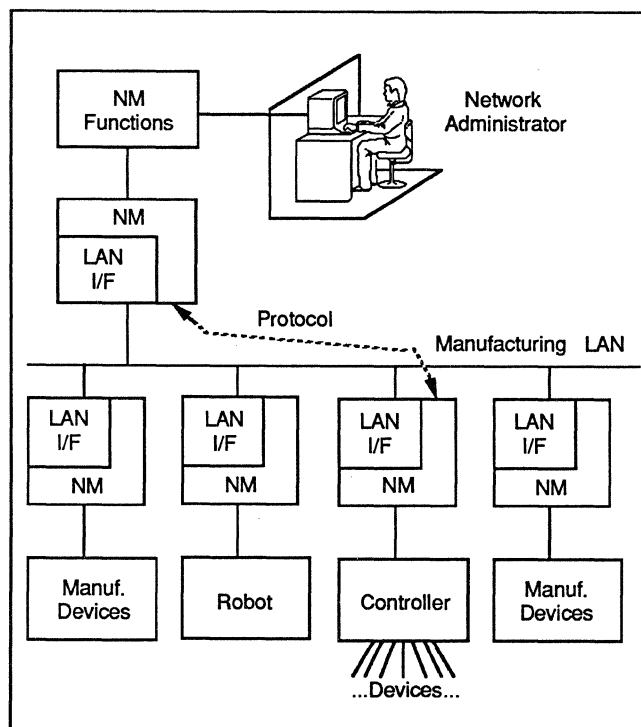


Figure 2. Manufacturing LAN with network management elements added.

Network Management Functions

The exact functions that network management should provide are still being defined by the standards groups. So far, the standards groups have agreed upon a set of basic functions that include configuration and name management, fault management, performance management, accounting management, and security management.

- **Configuration and Name Management** are mechanisms to determine and control the characteristics and state of a LAN, and to associate names with managed resources. Some of the services that configuration management provides include setting LAN parameters, initializing and terminating LAN resources, collecting data on LAN status for reports, and changing the LAN configuration. Some of the services that name management provides include naming the resources to be managed and managing name assignments.
- **Fault Management** includes mechanisms to detect, isolate, and recover from (or bypass) faults in the LAN. The way fault management is performed will depend on the LAN's application. For example, in some manufacturing applications, LAN down time is intolerable. In these cases, fault management should be proactive; that is, fault management should forecast probable faults and emphasize pre-

Network Management: A Manager's Perspective

ventive maintenance. In LANs where down time is not so catastrophic, fault management could be reactive, i.e., acting only in response to faults as they occur and emphasizing accurate diagnosis and rapid repair.

- **Performance Management** includes mechanisms to monitor and tune the LAN's performance, where performance is defined by user-set criteria. Some environments may need special performance metrics. The factory environment, for example, is typically hostile to communications equipment: noisy, dirty, and with wide temperature variants. Thus, instead of classic performance metrics that assume error-free operations, these LANs should use metrics more on the line of performability^{13,8}—a means of measuring performance and reliability in a unified manner.
- **Accounting Management** includes mechanisms for controlling and monitoring charges for the use of communications resources.
- **Security Management** includes mechanisms for ensuring network management operates properly and for protecting LAN resources.

Although accounting management and security management mechanisms have been defined, they have not yet been included in the MAP/TOP specifications, so we will not discuss them further in this report.

AN EXAMPLE NETWORK MANAGER

The standards efforts have defined the basic functions that network management should provide, but they have not defined how these functions should be applied to the management of a particular type of network or network application. That is, the standards provide a set of basic services, but they do not specify how to implement or use the services.

To provide a better idea of how these functions could be used, we give an example of a network manager application. This network manager is "ideal" in the sense that not all of these services may be readily implementable at this time.

Our example system supports the three network management functions specified in MAP/TOP 3.0: configuration management, fault management, and performance management. The services it can provide for each of these functions are summarized below.

Configuration Management

- **Defining the LAN topology:** the system provides tools to assist in setting up the network's design and initially configuring the LAN components. For each type of component, the system suggests how the component's configurable parameters should be set.
- **Displaying the LAN topology:** given the initial LAN design, the system generates a display that shows the location of and the operational status of each device in the LAN (active/inactive).
- **Reading current values:** the system enables the network administrator to request the current value of any of communications resources within the LAN components, for example, the number of messages that have been sent or received by a specific network device.
- **Setting values:** the system enables the network administrator to modify the value of the configurable parameters in the LAN components, such as the device's status (active/inactive), or certain protocol layer parameters (number of retransmissions at the transport layer).
- **Adding or deleting devices:** the system supports dynamic changes to the LAN's topology, such as adding or deleting components to/from the LAN.

Fault Management

- **Detecting and giving notices of faults:** the system gives the network administrator fast notice that a fault has occurred somewhere in the LAN. The system has an option to change modes from reactive fault management to proactive. In proactive mode, the system gives notices when it predicts that faults will occur.
- **Isolating faults:** the system assists the network administrator in determining where in the LAN topology the fault occurred, which device failed, and which portion of the device failed. This can be done in one of two modes: an automated mode where the system isolates the fault without intervention from the network administrator, or in assistant mode where the system provides suggestions to guide the network administrator in locating the fault.
- **Correcting or bypassing faults:** the system again has automatic and assistant modes. In automatic mode, the system corrects the fault or implements a bypass without intervention from the network administrator. In assistant mode, the system gives advice on how to correct or bypass the fault, but the network administrator must make the changes.

Network Management: A Manager's Perspective

Performance Management

- **Collecting statistics:** the system maintains a set of current and historical statistics on each network device in its domain.
- **Evaluating performance:** the system uses the statistics to calculate performance metrics, and evaluates these metrics against predefined user criteria for performance.
- **Reporting:** the system generates textual reports and graphic displays of the statistics and performance evaluation results. These reports can be periodic or ad hoc.
- **Tuning performance:** the system can do this in either of two modes: automatic or assistant. In automatic mode, the system dynamically monitors performance levels and adjusts parameters when they move out of range. In assistant mode; the system provides advice to the network administrator on how to tune performance, but the network administrator must enter the changes manually.
- **Cable monitors:** These products monitor the status of low-level devices on the cable plant, e.g., the amplifiers, and power supplies. They may also monitor the RF levels on the cable.
- **Modem monitors:** These products monitor the status of the modems used in the LAN.
- **Protocol monitors and analyzers:** These products passively collect information on the protocol transactions at one or more layers. They may calculate statistics, such as average number of specific types of network traffic, and generate reports. Analyzers may provide additional functions, such as identifying patterns in network traffic or capturing a specific type of traffic for closer examination.
- **Configuration support tools:** Some products are beginning to be available that provide off-line assistance in configuring the operating characteristics of LAN devices. Some products can also download software into LAN devices.

CURRENT STATUS

The example network management system provides many useful functions. The MAP/TOP User's Group has led the OSI community in defining the mechanisms needed to provide these functions. The MAP/TOP 3.0 specification includes an application layer network manager, a set of basic network management services, and a network management protocol. The other standards efforts, most notably the International Standards Organization (ISO), are progressing slowly in fully defining all of network management's functions. For further information on standards efforts, see Reference 2.

The ENTERPRISE Network Manager used at the ENTERPRISE Networking Event 1988 International is a prototype that provides a subset of the network management services defined in the MAP/TOP 3.0 specification.¹⁴ More specifically, the ENTERPRISE Network Manager provides configuration management functions which include monitoring, control, and support functions and performance management functions for generating graphical displays of network statistics and printing a textual summary of network devices.

The current products that are available for network management can be categorized as cable monitors, modem monitors, protocol monitors and analyzers, and configuration support tools. The product capabilities listed below are summarized from the most recent MAP/TOP Product Directory.³

While not yet commercially available, it is expected that products implementing the basic capabilities of the ENTERPRISE Network Manager will soon be released. In addition, several other companies are developing network management systems for open system networks. The main emphasis of these efforts has been towards managing telecommunications networks. The major efforts include AT&T's Unified Network Management Architecture (UNMA), IBM's NetView, and Codex's 9800 Network Management System.¹²

AT&T's UNMA is a virtual network management system. This system provides network management information at a customer site by linking the local site to a remote network control center via a protocol based on an Open Systems Interconnection (OSI) profile. Only a few vendors supply products conforming to this architecture. IBM's NetView is an integrated network management product.⁵ NetView's main component is its command facility which includes data collection, monitoring, and control functions. A number of vendors supply NetView-based products. The Codex 9800 Network Management System is based on the evolving OSI standards for network management and is currently limited to managing only Codex products.

Upcoming Capabilities

Some of the capabilities that are currently being planned include expanded network management functions, full ISO protocol support, nonstandard interface support, and multiple network manager support.

Network Management: A Manager's Perspective

- **Expanded network management functions:** The current LAN network management products focus on managing the lower protocol layers. To expand network management functions to cover all protocol layers, network management must be implemented in two areas: the network devices being managed and in the network manager application. The network devices that implement MAP/TOP 3.0 include network management functions. For network management applications, upcoming systems will provide tools that interact dynamically with the LAN to manage its components. These tools will provide flexible network management functions that can take the special needs of a LAN application into account. The MAP Configuration system (MAPCon)^{10,11} is an example of what will be forthcoming in this area. MAPCon is an expert system for configuration management that currently can statically configure LANs in an off-line mode. MAPCon's next versions will provide dynamic interaction with the LAN.
- **Protocol suites:** Network management products must support a full complement of the ISO protocols. These may be tailored to a specific platform, e.g., MAP or TOP.
- **Nonstandard interfaces:** In the real world, not all LAN interfaces use standard protocols. Still, the devices using these interfaces must be managed. MAP/TOP has acknowledged this problem, although they have not yet determined a solution. This is one area where creative vendor solutions may drive the standards efforts.
- **Multiple managers:** Organizations with a distributed management structure or geographically separated LANs may wish to use more than one network manager. This is another area that the standards efforts will be addressing in the future.

FUTURE OUTLOOK

What is the long-range outlook for network management? The standards efforts are progressing slowly. Most estimates are that the ISO work will not be completed until the mid-1990s. The MAP/TOP Users Group has taken the lead in defining some facilities sooner and may spur the standards organizations to progress faster in some areas.

Once the standards have been defined, building viable network management systems will still be a difficult job. The network management functions are interrelated and complex. The key to developing useful network management systems will be the use of technologies that include several from the domain of

artificial intelligence (AI). Some of the applicable technologies that may be used to develop systems for network management include the following:

- **Discrete event modeling:** permits detailed dynamic experiments with complex systems that can help answer "what-if" questions. For example, stochastic activity networks are particularly useful for modeling networks.^{13,9}
- **Statistical pattern recognition:** permits the monitoring of the quality of processes and products and makes corrections before the system drifts out of acceptable bounds.
- **Sensor fusion:** synthesizes the outputs of several sensors to derive parameters of interest that are not available from a single sensor.
- **Control theory:** includes a variety of techniques for distributed control.
- **Distributed artificial intelligence (AI):** includes techniques for distributed problem solving.
- **Diagnostic reasoning:** includes techniques for determining problem causes and solutions.
- **Game theory:** includes techniques for reaching optimal solutions for games with multiple players, which may be particularly useful for multiple manager systems.

How can these technologies be used to build network management systems? Some examples of network management applications using these technologies are as follows.

- **Configuration management:** Configuration management applications include giving advice on how to configure a wide range of LAN devices, dynamically adding and deleting devices to/from the LAN, and dynamically setting or modifying LAN device characteristics. As mentioned previously, one example of an existing expert system for configuration management is the MAP Configuration system (MAPCon). Applicable technologies include discrete event modeling, control theory, and distributed AI.
- **Performance management:** Performance management applications include dynamically evaluating the LAN's performance and identifying problem areas, suggesting key parameters to watch for while evaluating performance, suggesting parameter ranges when changes are needed, and dynamically tuning the LAN's performance. Applicable technol-

Network Management: A Manager's Perspective

ologies include discrete event modeling, statistical pattern recognition, sensor fusion, control theory, and game theory.

- **Fault management:** Fault management applications include detecting and isolating a specific type of fault, suggesting corrections or bypasses for a specific type of fault, and then adding more complex faults as experience is gained. Applicable technologies include diagnostic reasoning, distributed AI, discrete event modeling, statistical pattern recognition, and sensor fusion.

PLANNING ISSUES

Network management can provide a wide range of services. Determining which services are best for a particular organization and LAN application requires the careful consideration of a number of organizational and technical issues.

Organizational Issues

Below is a short list of issues that an organization should consider in planning its network management services. For more detail on organizational issues, refer to Reference 1.

- **Training:** Who in the organization needs training in network management? What types of training should these people receive? What types of training should net administrators receive?
- **Organizational commitment to OSI:** How committed is the organization to the concept of open systems? How many devices with nonstandard interfaces will have to be managed? What is the schedule, if any, for phasing out these devices?
- **Budget priorities:** What are the organization's investment goals? Is the emphasis on the long-range or the short-term? Network management and open systems are relatively high short-term expenses with long-term payoffs.
- **Management information needs:** Who will be permitted access to network management information? What types of information are needed? In what form? How often?
- **Security:** What level of security should be used to protect network management information and facilities?
- **Management architecture:** How should the network management structure relate to the organization's

management structure? Is a centralized or distributed structure more appropriate?

- **Management mode:** Should the network manager work in a proactive or reactive mode?

Technical Issues

Below is a short list of technical issues that should be considered before implementing network management.

- **Network management architecture:** Will the network management architecture be centralized or distributed? If multiple managers are needed, is the organization willing to do research on how to define the functions that have not yet been addressed by the standards groups? (See missing pieces, below.)
- **Missing pieces:** Two key areas have not yet been addressed by the standards groups: multiple managers and managing devices with nonstandard interfaces.
- **Multiple network managers:** The issues that must be answered include whether an additional protocol will be needed between managers, how responsibility will be divided between managers, and how information from multiple managers will be fused into a single, meaningful picture of LAN operations.
- **Nonstandard interfaces:** Until a wide range of products are available that support the standard layered protocols, many organizations will be faced with the issue of how to manage devices with nonstandard interfaces.
- **Application dependencies:** What characteristics of the organization's LAN application have special network management requirements? For example, LANs with real-time traffic need performance maintained within strict limits.
- **Expert systems:** Which expert system tools are applicable to meeting the organization's network management requirements? How critical are these requirements? Should the organization consider pushing technology development in these areas by funding research or conducting its own research?

REFERENCES

- ¹Fleischer, Mitchell, and Sinha, Manoj. Organizational Constraints on Network Management: A Systems Perspective. In *Proceedings of the ENTERPRISE Networking Event 1988 International*, 1988. Baltimore, Maryland.

Network Management: A Manager's Perspective

²Gottschalk, Gary. Management of Multivendor Networks. In *Proceedings of the ENTERPRISE Networking Event 1988 International*, 1988. Baltimore, Maryland.

³Industrial Technology Institute. *MAP/TOP Product Directory*. Society of Manufacturing Engineers, 1987.

⁴International Standards Organization, TC97/SC21/WG4. *Information Processing Systems—Open System Interconnection—Basic Reference Model Part 4—OSI Management Framework*. DIS, ISO, August, 1987.

⁵The LOCALNetter. Special Report: IBM and Network Management. *The LOCALNetter Newsletter* 7(10), October, 1987.

⁶MAP/TOP Users Group. *Manufacturing Automation Protocol Specification—Version 3.0*. Implementation Release, MAP/TOP Users Group, July, 1987.

⁷MAP/TOP Users Group. *Technical and Office Protocol Specification—Version 3.0*. Implementation Release, MAP/TOP Users Group, August, 1987.

⁸Meyer, J.F. On evaluating the performability of degradable computing systems. *IEEE Transactions on Computing* C-22:720-731, August, 1980.

⁹Meyer, J.F., Movaghar, A., and Sanders, W.H. Stochastic activity networks: Structure, behavior, and application. In *International Workshop on Timed Petri Nets*, Pages 106-115. July, 1985. Torino, Italy.

¹⁰Muralidhar, K.H., and Irish, B. MAPCon: An Expert System for Configuration of MAP Networks. *IEEE Journal in Selected Areas in Communications*, June, 1988.

¹¹Paranuk, H.V.D., Kindrick, J., and Muralidhar, K.H. MAPCon: An Expert System with Multiple Reasoning Objectives. In Kusiak, A. (editor), *Expert Systems Design and Management of Manufacturing Systems*. Taylor, Francis, to be published.

¹²Rosenberg, Robert. Are Users Up in the Air Over Network Management? *Data Communications*, December, 1987.

¹³Sanders, W.H., and Meyer, J.F. *METASAN: A Performability Evaluation Tool Based on Stochastic Activity Networks*. ITI 85-22.1, Industrial Technology Institute, June, 1986.

¹⁴Sparks, Steven; Behm, Jim; Joseph, Celia; Muralidhar, Kurudi. The ENTERPRISE Network Manager. In *Proceedings of the ENTERPRISE Networking Event 1988 International*, 1988. Baltimore, Maryland. □

Network Management: End-User Perspectives

This report will help you to:

- Identify the key end-user issues in network management implementation.
 - Establish the importance of planning for the network management system.
-
-

The role of network management has grown in strategic importance for today's end-user organizations, particularly those which are heavy users of information systems and telecommunications resources and which generally rely on information exchange as a key competitive weapon. Besides the competitive value of information exchange, there are several technical factors which are contributing to the strategic importance of network management:

- *Growing network complexity and sophistication:* High aggregate bandwidth channels create opportunities for consolidation, but also carry the risk of operational vulnerability. In addition, increasingly heterogeneous mixes of public and private network facilities (e.g., hybrid networks) are found in the corporate environment.
- *Escalating percentage of information and communication costs as they relate to total operating expenses:* There is an increasing need to identify cost/performance improvement opportunities. However, the complexity of today's networks makes such opportunities difficult to identify by intuition and experience alone.

The functions generally required by users in a network management system are shown in Figure 1. As shown, all functions are related in some way to both

This Datapro report is based on "Network Management: End-User Perspectives," by Martin Pyykkonen, Arthur D. Little, from *Telecommunications*, February 1989. © 1989 Horizon House-Microwave Inc. Reprinted by permission.

physical and logical network management. In practice, many of these functions are viewed as physical *or* logical, but not both. This is principally due to the different views of the world from telecommunications and MIS perspectives.

Since network management is still an evolving discipline in many user organizations, there are several dimensions of planning and implementation analysis to consider, including some fundamental issues as described below:

- centralized versus distributed network management;
- OSI-based network management;
- vendor packages versus homegrown systems;
- physical versus logical network management; and
- simple versus sophisticated network management.

CENTRALIZED VERSUS DISTRIBUTED

For a network size of any consequence, there must be some element of central control. This could be in the form of a central manager/operator with distributed network intelligence, or in the extreme, in the form of IBM's NetView master/slave configuration. Regardless, there are limits on the degree of distribution which can be effectively handled in network manage-

Network Management: End-User Perspectives

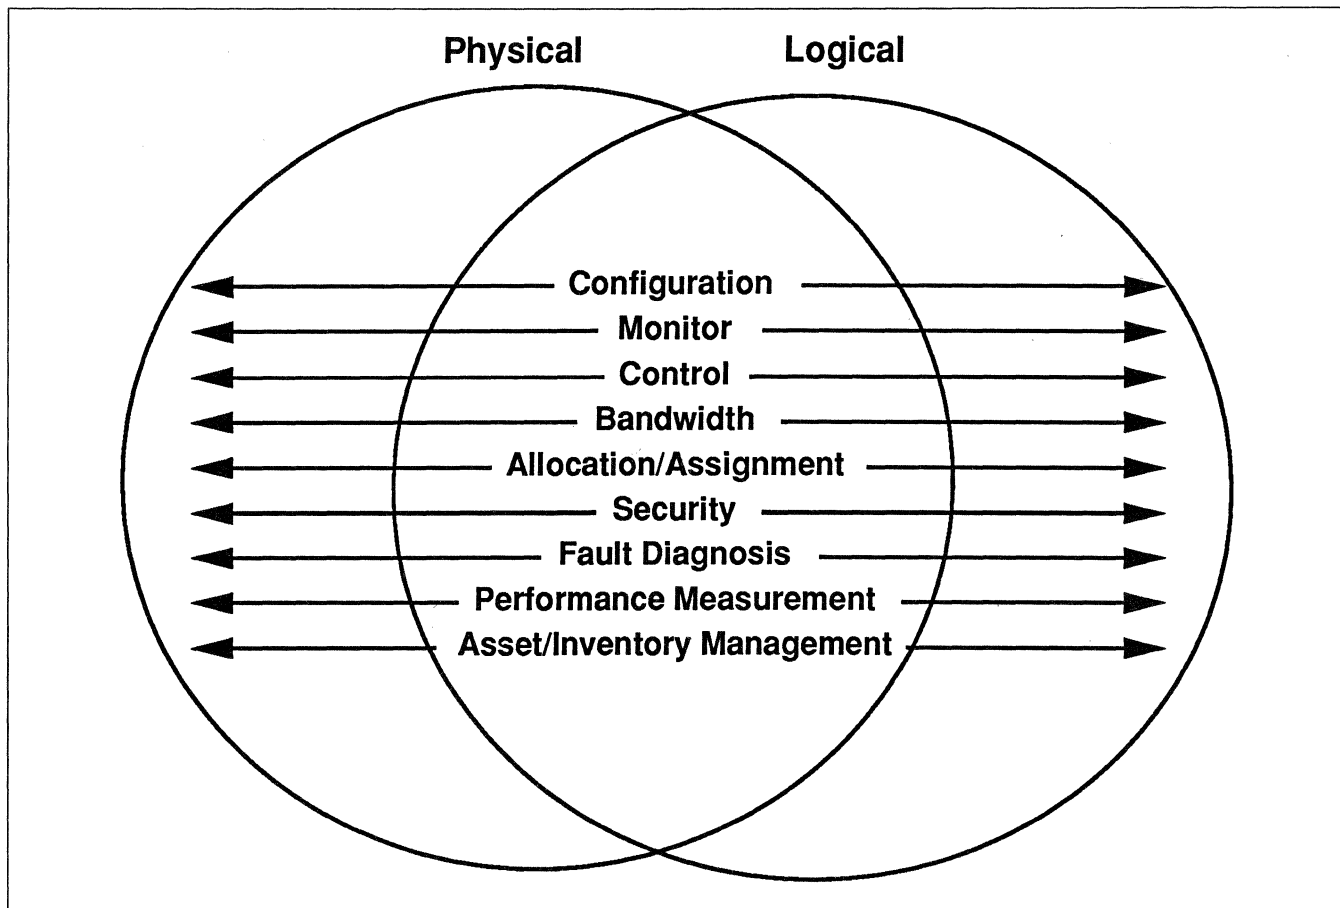


Figure 1. Most network management functions are related to a combination of physical and logical network operations.

ment, even more so than in the traditional debate over distributed versus centralized data processing in the MIS environment.

The industry has thus far been slanted toward "centralized" network management, based on critical reliability requirements in user networks and the marketing strength of IBM (NetView) in existing large and mainframe-based environments. DEC's recently announced statement of direction (EMA—Enterprise Management Architecture) will reopen the centralized-versus-distributed debate in some organizations. EMA will be an umbrella consolidation and an extension of DEC's network management strategy, namely, a DECnet-based integration of discrete packages.

OSI

An Open Systems Interconnection/Network Management (OSI/NM) group was recently established by eight founding members (Table 1). The goal of the OSI/NM group is to specify interface standards which are OSI-compatible and on which vendors can de-

velop their own network management products. The intended focus of OSI/NM is the future development of network management systems which conform to OSI. However, there is still the user need to deal with "today's" network management needs, regardless of OSI compatibility. Therefore, systems such as NetView have gained an early strong position with large end users. The debate centers on to what extent network management investments should be made in pre-OSI systems. Furthermore, the investment involves not only capital equipment, but also the less tangible costs of staff training and operations methodologies development.

Amdahl AT&T British Telecom Hewlett-Packard Northern Telecom STC PLC Telecom Canada Unisys

Table 1. OSI network management OSI/NM founding members.

Network Management: End-User Perspectives

VENDOR PACKAGES VERSUS HOMEGROWN SYSTEMS

Vendor-packaged solutions have been demonstrated to perform well operationally in many user environments, particularly when the functions are isolated to hardware diagnostics (e.g., line monitoring, response-time measurement, etc.). In order to meet the large number of user-specific needs in a comprehensive system, it is almost a standard requirement for the system to be at least partially customized. This is due to the uniqueness of user networks. There are probably no two networks which are sufficiently alike to be managed on an end-to-end basis by a standard off-the-shelf product.

PHYSICAL VERSUS LOGICAL NETWORK MANAGEMENT

Purely from a telecommunications perspective, a network does not experience a problem until there is clearly a physical failure—such as a line or circuit outage. At maximum sensitivity, there may be performance monitoring detection of soft error conditions which could later degenerate into an outage or failure. An individual user of the network, however, views the world from an “applications” perspective. A session failure will lead to an application outage even though the physical network continues to operate without incident. Herein lies the need to provide logical network session management if a claim of “end-to-end” network management is to be made.

SIMPLE VERSUS SOPHISTICATED NETWORK MANAGEMENT

Generally speaking, network management becomes more complicated as it becomes more logical (as op-

posed to physical) in nature, particularly due to applications software complexity versus rudimentary equipment monitoring in purely physical network management. The strongest factor involving sophistication/complexity is the extent that network management systems of *multiple* facilities or applications are integrated. Integration increases complexity and requires the user organization to be more sophisticated as well.

The implementation tools of a network management system can also be sophisticated, as in the case in which expert systems are utilized. Expert systems thus far have had mixed results—the most successful cases being where the project was kept simple and the goal was to eliminate the tediousness of mundane tasks, thereby reducing management labor costs. The most promising (and also most ambitious) longer-term potential for expert systems is analyzing the *interrelationship* of network events, particularly in the role of cause-and-effect diagnostic analysis.

It is difficult for most user organizations consistently to define their model of an ideal network management system. Individuals, of course, have a clearer view of the need and its resolution. For example, an MIS person will want complete management of a Systems Network Architecture (SNA) network (including session control, configure, activate, terminate, reactivate, etc.), whereas a telecommunications person will want circuit/line outage detection and fault isolation as primary objectives. For an organization to meld this range of needs requires accurate planning of the system objectives and implementation. And if multiple networks are to be managed in an integrated manner, the success of the network management system becomes even more dependent on accurate planning. □

Disciplines for Effective Network Management

This report will help you to:

- Design effective network management systems based on knowledge from multiple disciplines.
 - Analyze one approach to network analysis, user interfaces, network management protocols, and distributed architectures supporting an automated network management system.
-
-

Internetworks are managed with distributed administration and ownership and frequently contain many different types of equipment. Consequently, network management has become more important and more difficult. Effective network management requires knowledge from many disciplines including Communications, Network Analysis, Databases, Distributed Systems, Artificial Intelligence, and Human Factors. This report describes our research into these areas in support of the *Automated Network Management* (ANM) system which will be used by network operators, network analysts, and administrative personnel. This system is currently in development at BBN Laboratories Incorporated under contract to DARPA and CECOM and draws upon earlier work at BBN.

INTRODUCTION

This report describes our research in support of the *Automated Network Management* (ANM) system which will be used by network operators, network analysts, and administrative personnel. This system is currently in development at BBN Laboratories Incorporated under contract to DARPA and CECOM and draws upon earlier work at BBN.¹⁻⁷

This Datapro report is based on "Disciplines for Effective Network Management" by R.W. Callon, K. Corker, M. Nodine, J. Ong, M. Stillman, and J. Westcott, BBN Laboratories, Inc. © 1987 by IEEE. Reprinted with permission from the *Military Communications Conference 1987 Proceedings*, Washington, DC, October 12-22, 1987, pp. 19-26.

Distributed System Architecture

Figure 1 illustrates the distributed ANM system. Internetworks may contain many types of network components. Each Distributed Management Module (DMM) collects network management information and sends control commands to the components under its control. Client Processes communicate with the DMMs to provide the analyst with a user interface to ANM's data retrieval, data analysis, and control capabilities.

User Interfaces

We are designing a user interface that will intelligently assist a network operator or analyst by providing detailed and high-level information automatically or on demand. This user interface will present this information effectively and concisely via interactive text/graphics displays.

Data Analysis

Many analytical tools can be applied to assist the human user in determining network performance levels or detecting network problems. These tools include simple arithmetic calculations on raw data, statistical analyses, network analysis-specific algorithms, and AI-based reasoning.

Disciplines for Effective Network Management

Network Management Protocol

The network components, DMMs, and Client Processes will exchange network management information and control commands across the managed network using the Network Management Protocol (NMP) developed for ANM.

This report discusses several important aspects of the ANM system design. The second section, "Data and Situation Analysis," describes the forms of analysis that may be used by the ANM system. The third section, "User Interface," describes issues in effective user interface design and the approach that we have taken. The fourth section, "Network Management Protocol," describes the Network Management Protocol. The fifth section, "Network Management System Architecture," describes ANM's distributed architecture.

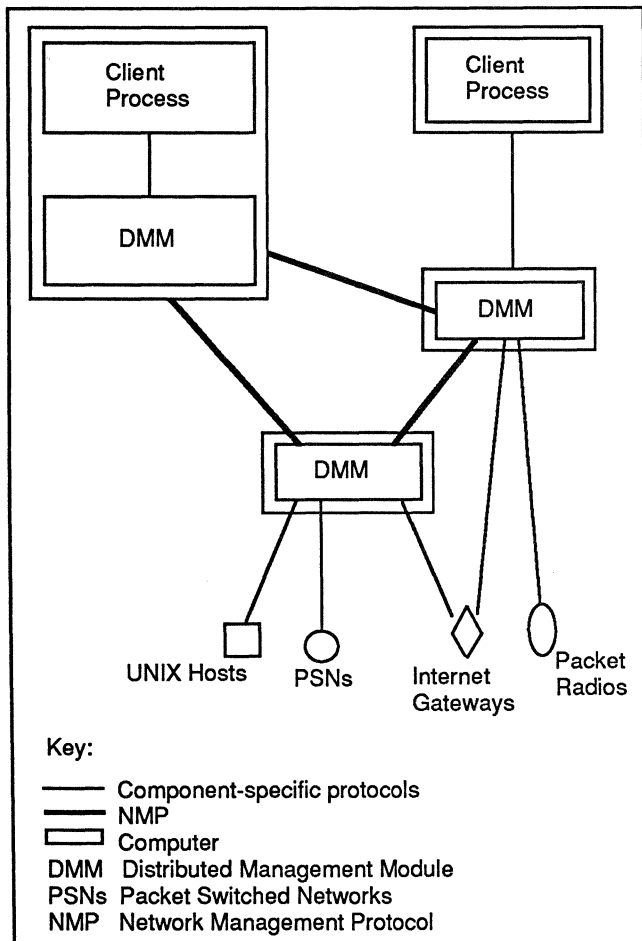


Figure 1. Automated Network Management (ANM) system architecture.

DATA AND SITUATION ANALYSIS

The ANM system will provide many analytic tools that convert the raw data received from the network components into meaningful and presentable assessments of the network's behavior. These tools may be classified by:

- The types of analysis.
- When the analyses are initiated.
- The users of each tool.

Types of Analyses

The network management system will answer many kinds of questions posed by the network analyst and will pose questions to itself when automatically detecting and diagnosing network problems. To answer these questions, the system will support many types of analysis ranging from simple arithmetic calculations to Artificial Intelligence-based reasoning.

The simplest analyses calculate a meaningful datum from two or more data values as received from the network components. For example, a packet switch may report the number of packets received and packets dropped, and the system may calculate the percentage of packets dropped. The system will also be able to answer database and statistics questions such as "Which packet switches have transmitted more than X packets?" or "What is the mean and variance for the queue lengths of all packet switches?" Many questions may require the application of network analysis algorithms that detect routing inconsistencies, traffic bottlenecks, or other abnormal states by using graph search, queuing theory, or other techniques. Other questions such as "Are there symptoms of congestion near packet switch Y?" may require the system to reason about concepts such as "congestion" at a high level using AI representation and programming techniques.

When Analyses Are Initiated

Collecting all network management data available from the components without purging will diminish CPU time, user excessive storage, and congest the monitored network. Frequent application of all available analyses will also cause the system to collect excessive data as well as consume inordinate processing resources. The network management system must therefore minimize routine data collection and analysis while ensuring that network problems are detected and diagnosed quickly and accurately.

Disciplines for Effective Network Management

Most automated analyses in our network management system will be triggered by easily detected events. Our system will collect regular updates sent by each network component and apply superficial analyses to detect small and simple patterns of data. If a detected pattern suggests a possible network problem, the system will then apply more computationally intensive analyses. If necessary, the system will send a query to a DMM or trace or echo packets through the network to obtain information not routinely collected.

It is sometimes difficult to design computationally-inexpensive triggers on small patterns of data that can reliably detect a possible network problem on an event-driven basis. In these situations, the system will execute diagnostic tests at regular time intervals.

The Users of Each Analytic Tool

Different network management personnel require their own analytic tools and apply these tools to various subsets of the available network management data.

Network operators are responsible for maintaining continuous operation of the network and must be informed of network failures as quickly as possible. They require access to the current status and recent history of the network to diagnose and correct the problems quickly.

Network analysts possess additional expertise which enables them to diagnose more subtle problems or improve algorithms applied by the network components. These people frequently require more comprehensive information and analysis about the network's current and recent status as well as historical information and cumulative statistics that may indicate chronic problems or suboptimal conditions.

Administrative personnel may require cumulative statistics related to network utilization and performance to insure that the network is optimally configured for the users' traffic loads and performance requirements.

USER INTERFACE

The user interface provides network operators, analysts, and administrative personnel with convenient access to the data collection, filtering, analysis, and presentation capabilities offered by the ANM system. In doing so, the user interface must support the functional requirements of network management and

complement the goals, expectations, and cognitive abilities of the user. This subsystem must:

- Determine which information is important,
- Display information effectively, and
- Support a dialogue with the user.

User Interface Objectives

Determining What Information Is Important. Network management databases are very volatile. In our system, for example, most of the database may be updated with new values approximately every fifteen minutes. Any single piece of data is potentially significant, but naively displaying every new value will overwhelm the user. The interface should distinguish important information from background information and abstract significant patterns of data that is distributed in time and in location.⁸ Information is important if it identifies the status of network components; suggests possibly incorrect network performance; helps confirm or reject a previously-generated hypotheses; or contributes to the user's development of an accurate model of the network. The user can request information from the system believed to be useful, and the network management system itself may receive or infer information which should be brought to the user's attention.

Effective Display of Information. The system must display answers to queries or data analysis requests so that the user can interpret and apply the information effectively to solve the problem at hand. The same information can be visually encoded in different ways. Each may support different network management goals and require different cognitive processing strategies by the user. For example, the operator may request the buffer utilization for all packet switches. This information might be displayed as a table containing the buffer utilization value next to each switch name for precise numerical analysis. Buffer utilization values could also be displayed in a histogram that quickly presents the distribution of values for comparative assessment. Finally, a graphics display could color-code an icon for each switch based on its buffer utilization to provide a quick "gestalt" status check and make important patterns of data more easily detectable.

Supporting a Dialogue. Answering isolated questions is, in itself, insufficient for effective communication between the user and the system. Detecting, diagnosing, and correcting network problems or inefficiencies requires the user and system to engage in an effective dialogue. This interaction can be improved by care-

Disciplines for Effective Network Management

fully designing the system's displays and commands to complement the user's mental model of network analysis.

Some dialogues are initiated by the user interface when the system automatically infers important events. Immediate display of this information is unacceptable: if the user is attending other tasks, he or she may miss the message. Instead, the system should indicate the presence of incoming messages and present summarized and detailed versions of each message on demand. Messages vary in type and urgency, so effective system management and presentation of these messages can help users prioritize the order of their display.

Other dialogues are directed by the operator. For example, to diagnose a network problem, the operator may ask a series of questions to support, reject, or suggest hypotheses about the problem's cause. An effective user interface provides an environment where the operator can easily relate an answer to previously presented information and use that answer to generate appropriate next questions.

Guided by the recent research in large-scale process monitoring and simulation,⁸⁻¹¹ we have assumed two basic models with which the user approaches network management. The first model is the *symptomatic approach* where the user makes a set of observations and checks them against known failure conditions or patterns of network behavior. Symptomatic management is valuable when an overall model or understanding of network behavior is available. The second model is the *topographic approach* where the user engages in a logical stepwise search to determine the network state, guided by functional, causal, and temporal relations among the network management data. The user interface must support information exchange in several ways to support these different mental models of the network management process.

Our Approach

We are structuring our user interface as a coordinated set of interactive text/graphics displays, called *views*; each provides the user with access to a subset of the system's capabilities. Dividing the user interface among a set of views has several advantages. First, each view's display can be optimized to display a specific type of information. Second, each view contains a particular set of displays and commands consistent with the way experts normally visualize the problem, so views can serve as an instructional aid to guide less-expert users.

Designing an effective view requires matching the cognitive aspects of the particular task with the abilities, knowledge, and *a priori* mental models and expectations of the user. This section overviews our view design guidelines, based upon human factors principles and research from other applications such as large process control systems.

Intuitive Encoding of Information via Visual Cues.

Many visual cues such as text, tables, and graphs can encode information. Using icons to represent physical and logical network entities enables the view designer to encode information using visual cues such as color, texture, shape, size, position, orientation, and blink frequency. People have expectations about the meanings of many visual cues (such as red for "alarm"), and views that conform to these expectations are easier to learn and interpret.

Consistent Encoding of Information via Visual Cues.

Confusion results when the user of a graphics-oriented system draws false associations between a visual cue and a meaning. We minimize this confusion by displaying legends that explain each cue's meaning and by adopting common conventions for visually encoding information. Consistently using the same coding conventions across views also avoids confusion and reduces the memory load required of the user.

Amount of Information per View. If too little information is displayed per view, the user may be forced to ask many questions to obtain the information actually desired. If too much is displayed, the user may be forced to perform excessive visual search and pattern matching to extract the desired information. We have created guidelines for the desired number of distinct visual cues presented at one time.

View Navigation. The *View Manager* portion of the user interface can organize the views and anticipate user needs by suggesting appropriate next views. Views can be organized in several ways. Some views can exploit natural hierarchy of the network. For example, a view may allow the user to mouse-select the icon representing the ARPANET to request a new view describing the network's switches and links. Views can also be organized by the type of user so that views which are inappropriate for a particular user are "hidden" to reduce confusion. Views can also be structured by the role each view takes in a larger task. For example, when the system notifies the user of a situation that exists in the network, the user may need additional information to understand the situation.

Disciplines for Effective Network Management

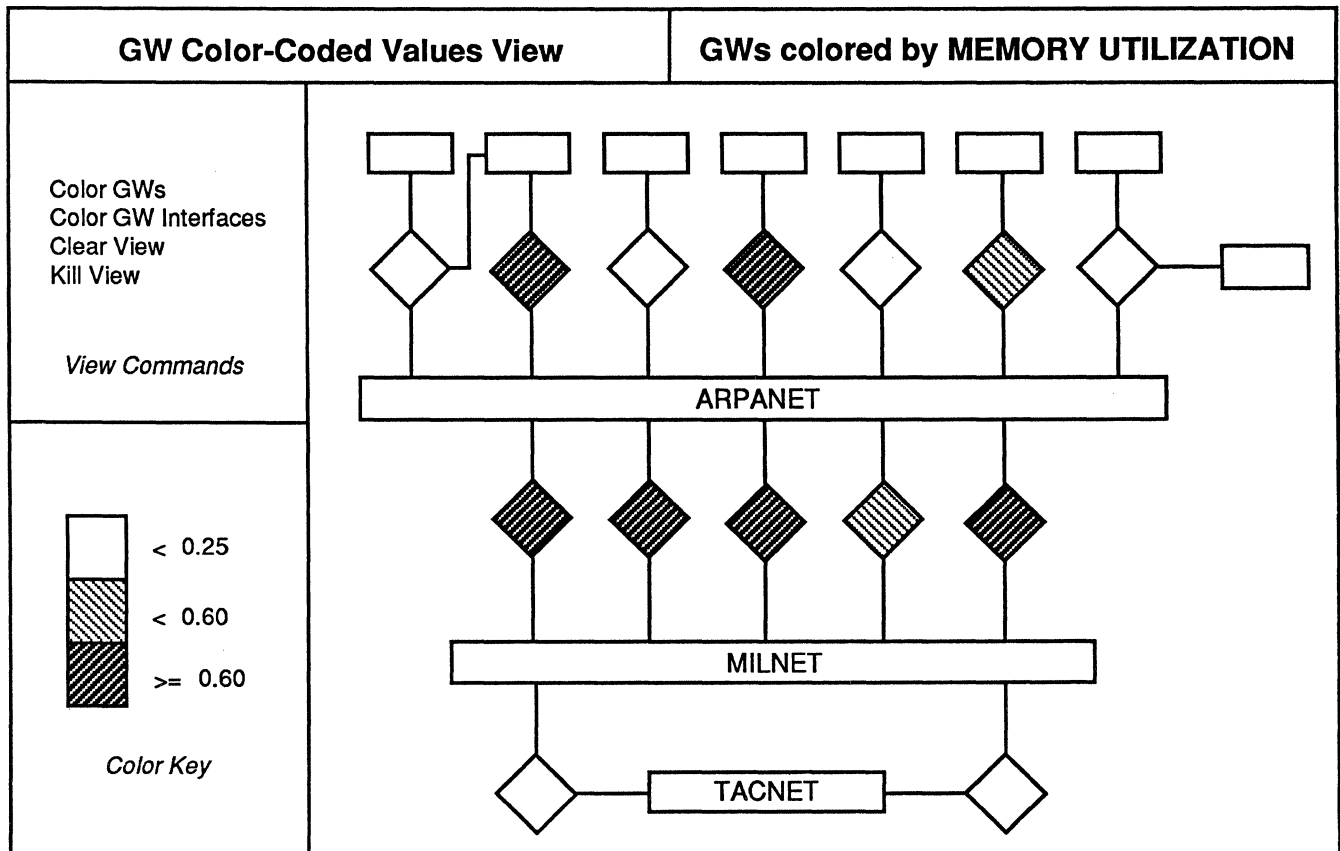


Figure 2. Example view of the Internet.

An Example

One of the views we have developed displays a diamond shaped icon for each gateway, a rectangular icon for each network, and a line-shaped icon for each gateway/network interface in an internet. This view displays the internetwork's topology by connecting the icon for each interface to the icons for its associated gateway and network. Figure 2 shows a simplified version of this view.

This view responds to two common and complementary types of queries. First, the user can supply an arithmetic expression or a function name to calculate a value for each network entity of interest and color-code the icons accordingly. Second, the view can print detailed information about a particular network entity when the mouse is clicked over the entity's icon.

This view facilitates dialogue by helping the user identify network entities requiring further investigation. By color-coding each entity's icon, the user can quickly identify significant patterns of data and network entities requiring further investigation. The user can then mouse-click on those icons for more detailed information. For example, in Figure 2, one can easily

detect that many of the gateways between the ARPANET and MILNET have a high memory utilization which may warrant closer investigation.

NETWORK MANAGEMENT PROTOCOL

DMMs and Client Processes exchange monitoring and information using the Network Management Protocol (NMP) developed for ANM. NMP consists of three subprotocols—the Query/Response Protocol (QRP), Network Component Control Protocol (NCCP), and the Management Session Protocol (MSP). This section discusses each of the subprotocols and additional protocol considerations.

Query Response Protocol

The *Query/Response Protocol* (QRP) is based on a simple query and response two-way handshake where a *requestor* such as a Client Process queries a *responder* such as a DMM for network management information about network components. Queries specify:

- The *range* of the network components queried,

Disciplines for Effective Network Management

- When the data is desired, and
- Which data is desired.

Ranges of Network Components. The *range* of a QRP query is the set of *items* for which data is desired. In most cases, the item is an actual physical component such as a gateway, packet switch, or host. It may also be a subcomponent, a device associated with a physical component such as a front-end processor, or even a logical structure such as an entire network. We have defined a simple but flexible format for defining ranges of queries and commands.

When the Data Is Desired. A query for information may request the information to be sent:

- When the query is received,
- At a specified time in the future,
- On a periodic basis at specified times in the future, or
- On an event driven basis, whenever the information changes during a specified time interval.

At times, it may be desirable for network components or DMMs to send information that has not been requested explicitly. For example, most network management centers receive fault reports and status information automatically. QRP supports this feature by allowing systems such as DMMs to send unsolicited responses.

Which Data Is Desired. A *clump* is a related set of data values about a network component. For example, one clump may contain throughput statistics for a packet switch. Many clumps may be specified for each network component type. A requestor can retrieve network management information by specifying the desired clump and set of components.

Requesting information by specifying a predefined clump is amenable to efficient implementation and is sufficient when the kinds of information desired by each query can be anticipated. If the network analyst or Client Process creates queries on the fly, a more flexible query language may be needed. We are considering extending QRP to encode queries expressed in a relational database query language.

Network Component Control Protocol

The *Network Component Control Protocol* (NCCP) supports control of network components. The NCCP

also uses a two-way handshake consisting of a command and an acknowledgment.

A control command specifies the control action desired, the network components to which the command applies, and the time at which the command should take place. Command actions currently supported by NCCP may set the value of parameters and may start, stop, or reset network activities within a device. Control actions may be performed at the time that they are received or at a specified time in the future.

Management Session Protocol

The QRP and NCCP use the services of the session-level¹² *Management Session Protocol* (MSP) which, in turn, uses a reliable transport protocol. Because access to the DMM database will be long-term, it will use a simple session-oriented protocol. Access to information within individual network components will be on an "occasionally get some information" basis and will use a transaction-based transport protocol and an essentially null session protocol. The transport protocol is not part of the NMP.

Protocol Considerations

Addressing. NMP addresses must be able to identify each DMM, Client Process, and network component in the internet. TCP/IP addresses are sufficient for most of the devices managed by our system. However, they cannot be used in environments such as OSI which use other addressing schemes or with network components such as packet switching nodes which are too "low down" in a network to own TCP/IP addresses. For these reasons, NMP addresses are specified by two fields: "address context" plus "address used within that context". Specific address contexts will be assigned for each addressing scheme such as "TCP/IP address", "OSI Transport plus Network Layer Address", "C30 packet switch address", and "DMM Identifier". Identifiers will be assigned to each DMM and Client Process.

Data Identify, Format, and Meaning. Different types of devices may report network management data using diverse representations. We make a distinction between a datum's *identity*, *format*, and *meaning*. For example, a particular table of packet counts could be identified as "standard gateway packet counts". The format of this table would allow a system to parse the table and determine the identity and data type for each field. The meaning would specify different aspects of the data to guide its interpretation and may be quite complex. The meaning of a field in a table of

Disciplines for Effective Network Management

packet counts may specify the type of packet counted, whether the field is zeroed each time it is read, how accurately the field is measured, etc.

The X.409 Presentation Transfer Syntax and Notation¹³ encodes the format and identity of fields, but it does not encode the meaning in the sense defined above. Nonetheless, we have found the flexibility and generality of X.409 to be useful, particularly where there is implied structure in the data being encoded.

Additional Protocol Mechanisms. Additional protocol mechanisms are needed for monitoring and control. For example, error reporting functions allow NMP entities to report on any errors detected in the operation of the NMP protocol itself. This is particularly useful for initial debugging of the network management system. Authentication and access control prevents unauthorized access to network management data and unauthorized control of network components.

NETWORK MANAGEMENT SYSTEM ARCHITECTURE

Architecture Overview

The network management system contains a set of cooperating Distributed Management Modules (DMMs) which collect and store information from the network components as well as provide an interface for controlling them. The network management information may be retrieved from the DMMs by Client Processes using the NMP protocol. Client Processes analyze and present the network management data to the user.

The Distributed Management Modules have four basic functions. First, they store monitoring data collected from network components in a database that can be queried by Client Processes and other DMMs. Second, they archive historical network management data for later retrieval and analysis. Third, they effect network control by translating and forwarding control commands from Client Processes to the network components. Finally, DMMs forward queries and control commands to Client Processes and other DMMs to support transparent, distributed network management.

Distributing the functionality of the network management system across multiple DMMs and their Client Processes has several advantages. First, it supports flexible allocation of DMMs and Client Processes to different logical or geographical subsets of the inter-

network. Second, careful deployment of DMMs and Client Processes can reduce traffic bottlenecks that may exist around a single centralized network control center. Third, additional DMMs and Client Processes may be added as needed to support the growing inter-network. Finally, distributed architectures reduce single points of failure and vulnerability to network partitions, making the system more survivable.

DMM System Architecture

The architecture of the Distributed Management Modules is shown in Figure 3. The submodules within each DMM can be grouped into those responsible for collecting and storing of information from network components and those responsible for retrieving this information for client processes.

Collection and Storage of Network Management Information

Data Gatherers/Controllers collect monitoring data from a set of network components and send them control commands. Currently, each type of network component may use a different monitoring protocol, so the DMM contains one Data Gatherer/Controller for each component type monitored. For example, the DMM maintains one Data Gatherer/Controller each for LSI-11 gateways, BBN C/30 packet switching nodes, and SUN workstations because they all use different monitoring protocols.

Monitoring data can be collected from the network components either automatically or in response to specific requests. In both cases, the Data Gatherer forwards the data to the Packet Recorder for storage in the Database and Data Archiver. If a query cannot be answered locally, the Data Gatherer forwards it to the appropriate remote DMM. This cooperation among DMMs supports distributed network management within the ANM architecture.

The **Derived Data Builder** performs simple computations on raw data to derive additional, meaningful data values. These derived values are stored in the Database and may be retrieved via QRP queries.

The **Data Gatherer Manager** forwards queries and control commands to the appropriate Data Gatherer/Controller or Derived Data Builder. If no such process exists, the Data Gatherer Manager initializes a new one.

The **Packet Recorder** records each message as it is received. Incoming messages to the Packet Recorder include monitoring packets, error responses, and con-

Disciplines for Effective Network Management

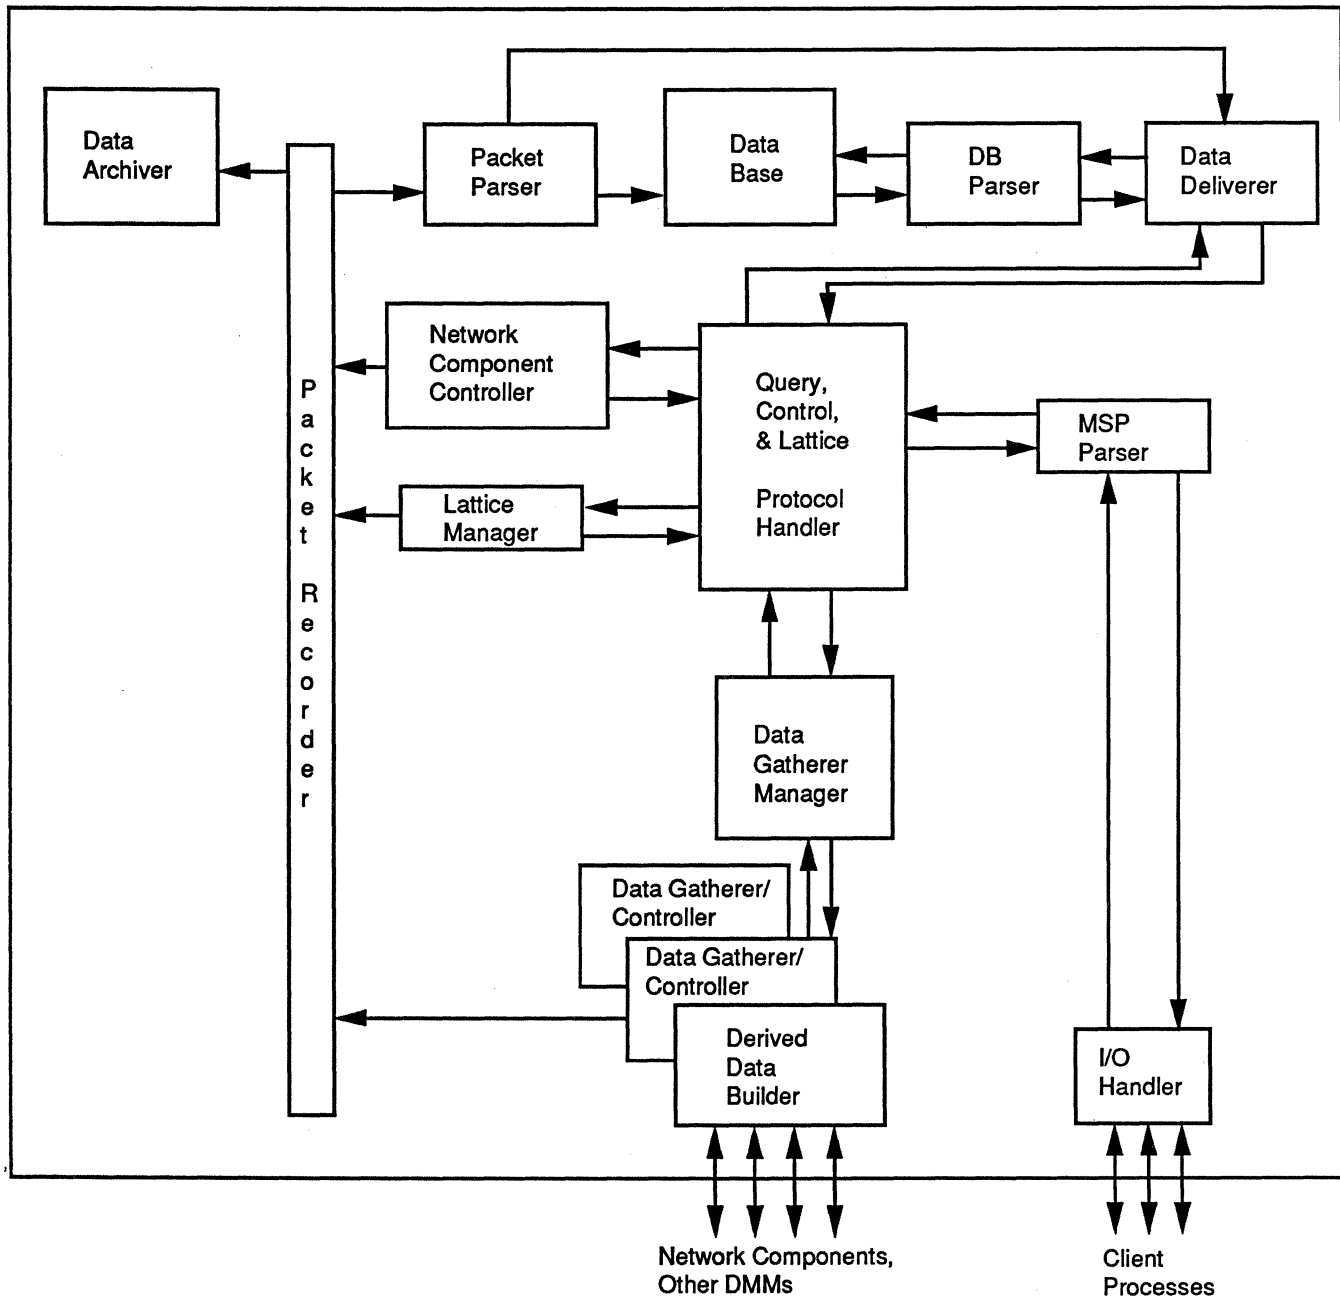


Figure 3. Distributed Management Module (DMM) system architecture.

Control responses which have been collected by the Data Gatherers. The Packet Recorder sends a copy of each message to the Data Archiver to be archived. It also forwards messages which are in sequence to the Packet Parser so that the information in it can be inserted into the Database.

The **Data Archiver** archives the incoming monitoring data. It also maintains indices into the archived messages so that they can be retrieved easily.

The **Packet Parser** parses incoming messages received from the packet recorder and places the information extracted from those messages into the Database. To do this, the packet parser must know the format of each type of monitoring message it can receive. This information is obtained from the Data Gatherers.

The **Database** provides quick access to the most recent network monitoring data received from each network entity. This information is updated by the

Disciplines for Effective Network Management

Packet Parser when new data is received. Data can be retrieved from the Database via the Database Parser.

Retrieval of Network Management Information

The **I/O Handler** manages transport layer connections between the DMM and the Client Processes.

The **MSP Parser** parses the session-level MSP protocol, a subset of the NMP protocol described in the fourth section, "Network Management Protocol." This module provides bookkeeping for multiple sessions and their correspondences with transport connections. If a Client Process terminates a transport connection while leaving the associated session active, the system will continue to collect monitoring information associated with the session so that the user can examine the collected data when the session is later reconnected.

The **Query/Control/Lattice Parser** supports the Query Response Protocol (QRP) and the Network Component Control Protocol (NCCP) functions. The QCL Parser forwards valid commands to the Network Component Controller and forwards commands requiring immediate effect on remote network components to the Data Gatherer Manager. The QCL Parser forwards all valid queries to the Data Deliverer, and if the information is not in the database, the QCL Parser invokes the Data Gatherer Manager to obtain the requested information. The QCL Parser rejects invalid queries and control commands and returns an error response to the originator of the query or command.

The **Lattice Manager** maintains information needed by the DMM to determine if a command or query can be satisfied by this DMM or if it must be forwarded to other DMMs. Specifically, it keeps track of the status of the DMMs and the network components managed by each DMM.

The **Network Component Controller** parses control commands and performs access control and control arbitration. It prevents two different users from executing conflicting control commands and allows Client Processes to request commands to be performed at some time in the future.

The **Data Deliverer** receives queries from the Query Protocol Handler and keeps track of each query's status. It forwards the query to the Database Parser at the appropriate time(s). When it receives a response from the Database Parser, it forwards it to the Query Handler. Redundant queries are filtered by the Data

Deliverer, and multiple copies of the response are sent to all users requesting the information.

The **Database Parser** retrieves information from the Database at the request of the Data Deliverer. It knows the structure of the Database, so it can access the requested data efficiently. Currently, the Database Parser retrieves and returns predefined data sets corresponding to the NMP *clumps* described in the fourth section. In the future, the Database Parser also may be able to interpret queries expressed in a relational query language such as SQL.

CONCLUSION

Building an effective network management system requires the application of many disciplines. This report has described our approach to network analysis, user interfaces, network management protocols, and distributed architectures in support of the ANM system.

REFERENCES

- ¹D. Davis, "Managing Computer Communications," *High Technology*, March 1987.
- ²J. Westcott, J. Burruss, and V. Begg, "Automated Network Management," *Proceedings of IEEE Infocom 85*, March 1985.
- ³C. Topolcic and J. Kaiser, "The SATNET Monitoring System," *IEEE Military Communications Conference*, October 1985.
- ⁴S. Blumenthal, "Experiences in Building a Wideband Packet Satellite Network," *Networks*, 1985.
- ⁵S. Bernstein and J. Herman, "NU: A Network Monitoring, Control, and Management System," *IEEE International Conference on Communications*, June 1983.
- ⁶P.W. Cudhea, D.A. McNeil, and D.L. Mills, "SATNET Operations," *AIAA 9th Communications Satellite Systems Conference*, March 1982.
- ⁷P. Santos, B. Chalstrom, J. Linn, and J. Herman, *Architecture of a Network Monitoring, Control, and Management System*, Bolt Beranek and Newman Incorporated, 1980.
- ⁸D. Woods, "Human Factors Challenges in Process Control," pp. 1720-1770, in G. Salvendy (Ed.), *Handbook of Human Factors*, J. Wiley and Sons, 1987.
- ⁹J. Rasmussen, "Models of Mental Strategies in Process Plant Diagnosis," in J. Rasmussen and W.B. Rouse (Eds.), *Human Detection and Diagnosis of System Failures*, 1981.
- ¹⁰J. Rasmussen, "Strategies for State Identification and Diagnosis in Supervisory Control Tasks and Design of Computer Based Support Systems," in W.B. (Ed.), *Man-Machine Research*, 1984.
- ¹¹K. Corker, L. Davis, B. Papazian, and R. Pew, *Development of an Advanced Task Analysis Methodology for the Army-NASA Aircrew/Aircraft Integration Program*, BBN Labs report number 6431, December 1986.
- ¹²International Standards Organization, *Information Processing Systems—Open Systems Interconnection—Basic Reference Model*, ISO 7498.
- ¹³*Message Handling Systems: Presentation Transfer Syntax and Notation*, CCITT Draft Recommendation X.409, 1983. □

Evolving Market Opportunities In Network Management

This report will help you to:

- Confront the trends of increasing complexity and escalating costs in monitoring and controlling network operations.
 - Sort out the plethora of products that claim “network management” capabilities and understand what each product can and cannot do for your network.
 - Exploit emerging network management market opportunities to your best advantage.
-
-

Network Management is one of the least known areas in telecommunications and promises to be the fastest growing. This rapidly evolving market opportunity was borne from:

- the explosive technological developments in communications,
- the proliferation of vendors in the postdivestiture environment,
- the increasing demand for private networks, and
- the necessity for 100 percent communications availability to avoid revenue loss.

Effective network management is critical in the private network world, because without controls, corporations will waste hundreds of millions of dollars annually. The market for independent product and service offerings in network management and control diagnostics (NMCD) is conservatively estimated at \$650 million in 1988, with an anticipated growth rate of over 25 percent in the next 5 years. This report will

outline the reasons for the development of this market, describe the technology, and highlight key players in the areas of product and service offerings and network integration.

WHAT HAS LED TO THE EVOLUTION OF THIS MARKET?

Prior to the explosive technological growth in communications, a handful of vendors were responsible for installing and maintaining all voice, data and other services in a user's communication system. The user had every reason to remain confident that any piece of equipment or service he/she wanted, whether voice or data, could be integrated into the network by these vendors. However, the situation is now dramatically different and the need to manage complex networks is creating new opportunities of particular interest for investors in communications.

In the last few years, there has been a proliferation of data processing, computer, and modem terminal equipment companies as well as voice carriers, data carriers, and special service networks with all of these vendors fighting for account control. Since 1983, more than 300 new vendors have entered the telecommunications industry and more are on the way as new opportunities emerge. The situation is further exacerbated by the fact that each will give only mini-

This Datapro report is based on *Evolving Market Opportunities in Network Management* by Clifford H. Higgerson, Roland A. Van der Meer, and Laurie A. Kappe, Communications Ventures, Inc. © 1986, 1989, Hambrecht & Quist Venture Partners. Reprinted with permission.

Evolving Market Opportunities In Network Management

mal support to interface their gear with other vendors' equipment or with the network carriers. In addition, many network managers report a general lack of service attitude on the part of most vendors and difficulties in pinpointing responsibility. As a result, users are no longer guaranteed compatibility and interoperability.

At the same time as the networks grow more complex, voice and data communication usage is growing at an average rate of 20 percent per year, while the associated management costs are climbing at a rate of 35 percent per year. The user is compelled, both economically and technologically, to take control of the network. According to the Yankee Group, network costs are second only to salaries as an organizational expense with an average monthly bill exceeding \$5 million for each of the world's largest corporations with private networks. Both large corporate networks and small networks are affected by this problem. Large networks need systems, software tools, and techniques to be able to manage the increasing complexity of their ever growing network. Small business communications networks, which may consist of a dozen tielines and leased circuits along with associated switching and multiplexing gear, require 99.9 percent uptime and need the ability to ensure this service.

The burden of management and control is falling heavily on in-house network managers. They have the impossible task of reducing costs and improving performance as their world doubles in complexity. Effective network management has become the "hot button" in the private network world today. Those companies which offer products or service to ease that burden will be increasingly in demand with an estimated market growth rate of 30 to 35 percent. Communications equipment vendors who must sell into the networks will not be able to market their products without integrating network management capabilities. We estimate that these vendors will be able to attribute 5 percent to 40 percent of their revenues to this function over the next few years.

WHAT IS THE NETWORK?

In the simplest terms, a communications network is that entity that exists beyond the phone, PC terminal, or computer. For the vast majority of users, their only concern with this network is that the connection is made properly and problem free. The telecommunications manager, data communications manager, or network operator views the network somewhat differently. He/she must understand the functional concepts of the network as well as its technical detail. It is

his/her responsibility to make sure the end user never has to think about the network.

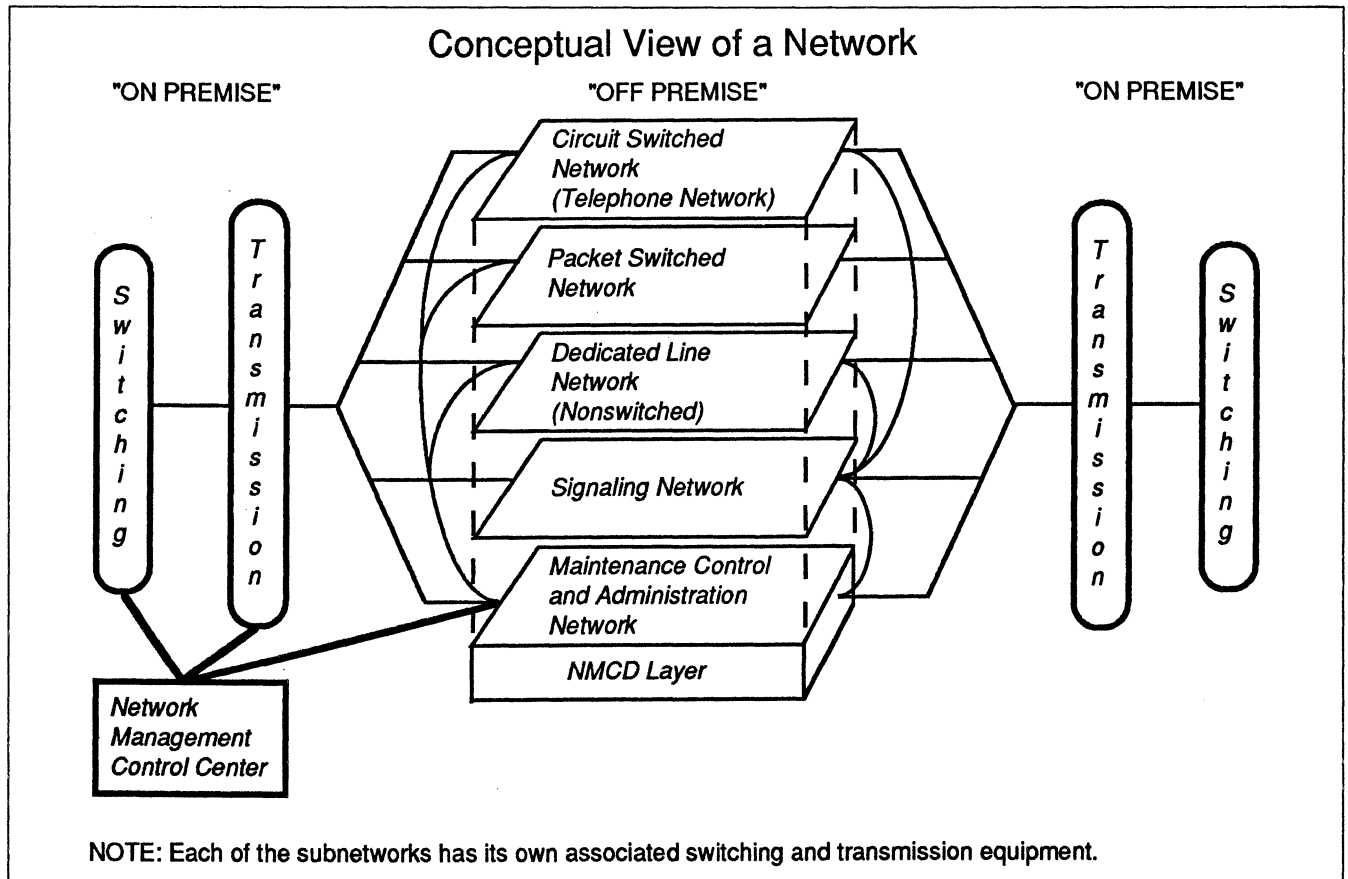
From the manager's viewpoint, the network is divided into "on-premise" and "off-premise" equipment. The "on-premise" equipment consists of both switching and transmission gear. Switching gear includes PABXs, key systems, data switches, and local area networks. The transmission gear consists of modems, multiplexers, digital cross connects, front end processors, and channel banks. The "off-premise facilities" consist of transmission facilities which may be owned by the user, leased from facilities carriers (such as MCI or the RBHCs), or virtual, which is the leasing of a logical point to point route from a facilities carrier or finally, used on demand, such as normal dial-up service. These transmission facilities can utilize one or more of the following media: microwave, satellite, fiber and wire cable. The off-premise facilities also include the associated tandem switching and multiplexing equipment. Figure 1 shows a conceptual view of the various subnetworks within the network. This figure illustrates the interrelationships between network management functions and the subnetworks.

The NMCD Layer of the Network

This report focuses on one aspect of the network, i.e., the Network Management Control and Diagnostic (NMCD) layer of the network. This layer includes the network and control functions, performance monitoring, troubleshooting and administration. The NMCD layer of the network is that aspect of the process that even an informed end user never worries about or even knows exists. However, the fact is that this layer is probably the most critical to the operation of the network as well as the most difficult to install and operate. Until recently, few people have appreciated the key role this layer plays in the effective performance of a network.

The primary trend affecting the operation of a network is the increased need for integration of systems and services. In the current environment, costs of equipment have leveled or declined, while the costs associated with monitoring and controlling these vast new utilities are escalating—and without proper monitoring, service levels will decline. In addition, without proper management controls, costs increase while system utilization may suffer from too much or too little equipment. The integration of voice and data networks, with required controls and data, is a relatively new demand. In some of the larger networks there could be as many as five to ten separate networks within a network to accomplish all of these functions. The multiplicity of these control networks

Evolving Market Opportunities In Network Management



Source: Communications Ventures.

Figure 1. Conceptual view of a network.

has been the result of adding them as the need evolved. There is a strong push to integrate the separate subnetworks in order to reduce the cost of the redundant facilities, functions and operations.

The emphasis of a network management system will change as the size of the network changes. For the small networks uptime is critical; uptime is basically a measurement of availability of the circuits in the network. For the medium size networks efficiency becomes important; efficiency is a measure of utilization as well as response time of the network. In the larger networks capacity monitoring becomes important; capacity monitoring is essentially the maintaining of data bases of usage statistics, and running simulations to enable the network manager to plan for growth.

For the purpose of this report, network management is broken into four components: network control, administration, performance monitoring and network troubleshooting. Table 1 summarizes each of the functions.

Network management tools today must be designed with effective human interfaces, color graphics, and windowing techniques, as well as high-level command languages. These tools allow the network manager to examine the network from a very broad overview, to rapidly focus on trouble spots, and to examine the network in a detailed piece by piece perspective. A new generation of proactive systems will slowly be incorporated into network management systems. This equipment will monitor, alarm, and restore trouble spots before any loss of service occurs. These functions will relieve the network managers of their hour-to-hour laborious tasks and allow them to focus on new and improved services. The network management education process is slowly evolving; there are currently no effective conforming standards. This means, in essence, that experience will be the only guide for many network managers.

THE EVOLVING MARKET OPPORTUNITY

The market opportunity for network management systems is being driven by the changing communications environment and the increasingly more com-

Evolving Market Opportunities In Network Management

Network Control Function <ul style="list-style-type: none">—Service type, service priority, mode of connection, bandwidth establishment of connection—Access security—Addressing—Routing strategy—Flow control—Signaling	Performance Monitoring <ul style="list-style-type: none">—Uptime efficiency performance analyses—Access availability—Call setup time—Network delay—Blocking probability—Switch performance—BIT error rate—Signal/noise ratio—Traffic monitoring (hold time, peak call time, traffic type)
Network Management and Administration <ul style="list-style-type: none">—Data base maintenance—Network equipment/inventory—Problem management—Installation management—Change management—Traffic and performance data base and reports—Traffic reports—Directory—Billing—Financial management	Troubleshooting <ul style="list-style-type: none">—Alarm monitoring—Transmission testing—Fault isolation—Diagnostics—Diagnosis—Fault rectification/restoral

Table 1. Network Management Control and Diagnostic (NMCD) breakdown.

plex technology. The absence of “hand holding” by the old AT&T has created the need both for self-reliance and for in-house managers to ensure effective control and cost containment. Five primary reasons for the development and growth of this market are summarized below:

- The rapid increase in both the numbers and costs of private networks; according to International Resource Development, Inc., over \$14 billion was spent in 1988 on U.S. private telecommunication networks and it is estimated that over \$40 billion will be spent in 1995.
- The increasing number of vendors in the market has made it imperative that equipment be made compatible with operations systems to facilitate network efficiency.
- The network operators, with increased performance demands and severe budget pressures, are continually looking for ways to improve efficiencies.
- The severe lack of qualified technical personnel, which has forced network operators to centralize and potentially automate network control.
- Increased reliance on data communications mandates accurate and secure transmission.

In order to identify the opportunities in this evolving market that telecommunication companies can and are exploring, we have divided the market into three categories: **product offering**, **service offering**, and **integration into existing products**.

Product Offering

As a standalone product offering, network management products can encompass one or all of the network management functions. From a technical point of view, the upper level network manager tools are the most interesting, i.e., a system that enables an operator to ensure the network is up and running and performing well. These tools attempt to integrate an entire system so the network can be monitored and operated in a centralized manner. We estimate the market for these systems to be over \$50 million today and growing at over 30 percent.

The key to network management product offerings is the ability to integrate many network tools into one control system. Integration is a necessity in today's environment; without this feature a network could conceivably have twenty different control centers. As integrated control systems evolve and network management becomes more understood, we will see expert systems taking over some of the decision making process in network flow control and restoral. Over the next few years we anticipate more robust system control from these types of packages.

A related category of product offerings is the administration and management of a network. This can include installation management, problem resolution, restoral management, change and financial management as well as some other areas. There is an increasingly critical need for such systems; network managers can spend 80 percent of their time keeping track of these aspects of the network, and a dedicated automated system could conceivably reduce operations costs by as much as 30 percent. In our review of

Evolving Market Opportunities In Network Management

current offerings we found four such systems, including one from IBM; however, the most advanced was the offering from Peregrine Systems.

The final category of NMCD product offerings are test, diagnostic, patching and switching equipment. These systems are primarily used by the technical counterparts to the end network. In larger networks, they operate "underneath" those systems. The companies in this area tend to be "hardware oriented" and focus on selling test sets, switchable patch panels, and remote diagnostics systems as well as some form of information management systems for the network. There is an overlap in the functionality of these type of systems and the performance monitoring and control systems previously discussed; however, in the larger networks they operate in unison.

The overall market size for NMCD product offerings is very difficult to derive primarily because there are no distinct boundaries on how much communications test gear should be included in these estimates. Our survey of industry experts has led us to conclude that the *minimum* market size for the NMCD products is \$650 million and it's growing at a rate of over 20 percent.

Service Offering

Network management as a service offering is most definitely not a new area—it is as old as the telephone network itself. The difference now is that the servicer is no longer taken for granted. Communications is a business where economies of scale dominate; very few medium- to small-size businesses can afford the time, cost, and expertise necessary to run a network. Since divestiture, the difficulties have compounded and we now see network management service centers turned into high-margin profit centers. The technical assistance centers (TAC), i.e., switching or multiplexing vendors that have been in business for decades servicing the telephone companies and long-distance carriers, have now entered the private network market.

The need for network management in the private area has now grown to the level that significant profits can be derived, thereby leading to a variety of independent service offerings. One unique business in this area is a nationwide offering from PacTel Spectrum Services, a company acquired by IBM in 1988. This service includes remote diagnostics, degradation monitoring, and coordination of rapid restoral. IBM PacTel Spectrum Service's offering is geared to clients with geographically dispersed networks comprised of multivendor networks that are critical to the success

MARKET SECTOR	1988	5-Year Growth Rate
Product Offering		
Performance Monitoring and Network Control	\$ 50M	30%
Network Management and Administration (NetView, ACCUMASTER, Net/Master, etc.)	\$220M	35%
Test, Diagnostics, Patching and Switching	\$380M	5%
Service Offering	\$ 95M	30%

Source: Communications Ventures.

Table 2. Market estimates for network management and control.

of their business. Communications equipment companies are also moving to exploit the need for servicing private networks.

RBHCs are taking advantage of the confusion resulting from the private network proliferation and have begun to offer services through divisions or subsidiaries. For example, NYNEX is currently looking for trail sites for its Intellihub virtual private network service, and Ameritech is also offering limited network management services. In addition, AT&T, MCI, and US Sprint offer software defined network services to allow a customer to directly define and control services as an inducement to use the carrier's network as opposed to a private system. These virtual private networking offerings allow customers to use as much or as little network capacity as they need on a demand basis. If these services are successful, the growth rate of the independent network management industry could be somewhat less than our projections.

Despite the involvement by the carriers and RBHCs, we still expect to see a few communications companies offering integrated network management services over the next five years. The total industry market size for network management services could reach \$200 million by 1990. The opportunity is tremendous for those who can find the right formula of service and marketing.

Integration into Existing Product Offerings

This area of the network management industry is not just an opportunity for communications companies, but a necessity for ongoing success in the marketplace. It does not matter whether the product is a multiplexer, modem, switch, channel bank, or digital cross connect—it will become almost impossible to

Evolving Market Opportunities In Network Management

successfully sell without some level of network management capability built into the system. Clearly, the more sophisticated and capable the controls, the better are the chances to sell the system. In addition, those vendors who include flexible interfaces for compatibility and effective human interfaces will have a distinct advantage in the increasingly competitive market-place.

Spare ports on communications equipment, which are crucial to tying into a network management system, are becoming fundamental features for new products. Network managers are demanding that communications devices be able to relay back their

own status and the traffic being handled. Many modem and switching vendors are now offering varying degrees of network management systems. This can create a significant problem for network operators, however, as there is often a different controller, PC AT, or other type of equipment for each vendor's set of equipment. Interoperability and compatibility is thus becoming a critical issue—creating the need to form standards committees, although it will be some time before the effect of this effort will be felt. It is our contention that successful communications companies will find that anywhere from 5 percent to 40 percent of their sales will be network management related. □

Network Management Functions: Telecommunications Hardware

This report will help you to:

- Identify key network management functions of network hardware from the front-end processor on out.
 - Gain familiarity with the nomenclature of emerging OSI management standards.
 - Evaluate the limitations of the five functional areas within the OSI Management model.
-
-

Network management functions are those which support the framework of network management systems (NMSs). Network management functions must be performed by the NMS and/or the network management staff to ensure efficient and effective operations of the network.

All network management functions are related in some way to both physical and logical network management.

Logical network management focuses upon the logical connection between the user and the ultimate network destination, which is generally a software application. The two principal components of logical network management are session awareness and traffic flow visibility. (See Figure 1.)

Physical network management involves failure notification, problem detection and identification, problem isolation, and problem resolution of physical entities. These entities include the circuits, lines, modems, multiplexers, switches, front-end processors, and other hardware components. (See Figure 2.)

Generally speaking, network management becomes more complicated as it becomes more logical in nature. This is particularly true when one compares applications software complexity with the basic equipment monitoring achieved with purely physical network monitoring and control.

The traditional separation of logical and physical network management has developed for three reasons:

- There has been a constant migration from simple configurations to complex ones.
- Most vendors provide only physical management products.
- The physical network management perspective is reinforced by common carriers selling transport services only.

There is a growing need for powerful relational network management databases for integrating information from both logical and physical aspects, particularly with regard to inventory, configuration, status, and performance history.

Portions of this report were contributed by Daniel Minoli. Mr. Minoli is an adjunct assistant professor at New York University's Information Technology Institute, as well as a full-time data communications researcher and strategic planner.

Network Management Functions: Telecommunications Hardware

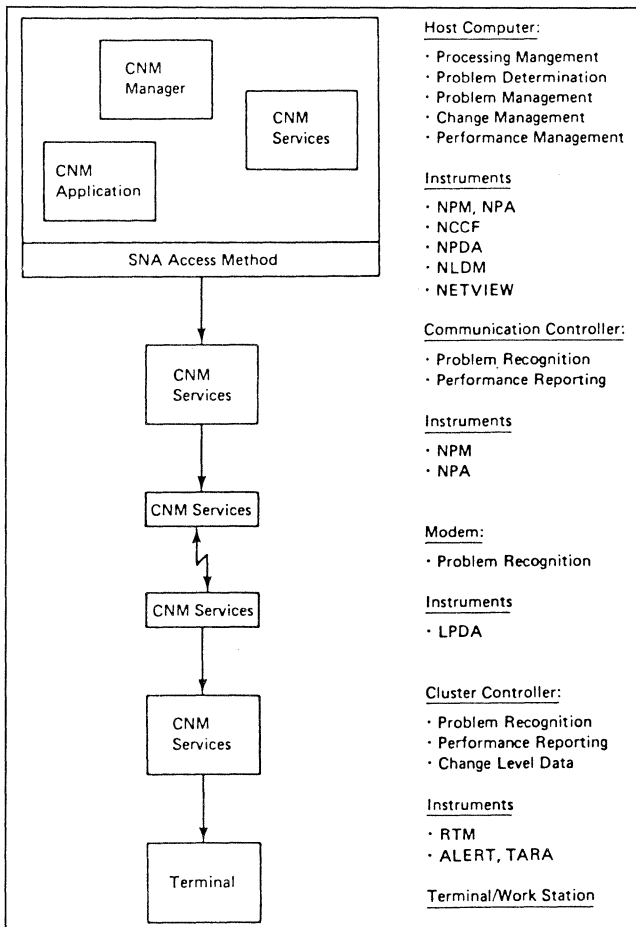


Figure 1. The typical architecture of a logical network management system.

NETWORK MANAGEMENT REQUIREMENTS

In general, network management requirements may be summarized in terms of these goals: enhancing the quality of service, increasing productivity, reducing complexity, and providing control.

To ensure efficient network management, the following key factors must be considered:

- **Network management functions**—the specific functions that must be performed by either the network management system or the network management organization. For telecommunications hardware, these requirements include, but are not limited to fault, configuration, performance, security, and accounting management.
- **Network management architecture**—the framework of network management functions and their placement within the network hierarchy.

- **Network management instrumentation**—the set of tools for collecting, compressing, and processing information and instruments for predicting future service levels and resource usage of telecommunications hardware.
- **Network management software**—the controlling software and the inventory/configuration database which stores information on all related network devices and components.
- **Organizational components**—the organizational structure put in place to carry out network management functions.
- **Organizational approach**—factors such as personnel skills and education that are critical to the successful long-term operation of human resources in management.

This report examines the first factor, network management (NM) functions of telecommunications hardware. There are at least as many definitions of NM products as there are vendors that provide NM products. Therefore, Datapro has chosen the five functional categories defined by OSI as a starting point for structuring its coverage of NM functions. Datapro also recognizes certain critical network management functions not included in the OSI classification; these categories are discussed in the following section.

OSI MANAGEMENT FUNCTIONS

OSI Management standards define five functional areas pertinent to management activities. Both vendors and users are now beginning to apply the OSI nomenclature to describe both OSI and non-OSI network management environments. The five OSI Management functional areas are: Fault Management, Performance Management, Configuration Management, Security Management, and Accounting Management. (See Table 1.) These areas focus on management of telecommunications hardware, from the front-end processor on out through the network.

Fault Management—encompasses fault detection, isolation, and the correction of abnormal operation. Faults cause communications systems to fall short of their operational objectives. Faults may be persistent or transient and manifest themselves as *events* to the network management system (NMS). NMSs with error detection capabilities can recognize faults. Fault Management functions include:

- Maintaining and examining error logs.

Network Management Functions: Telecommunications Hardware

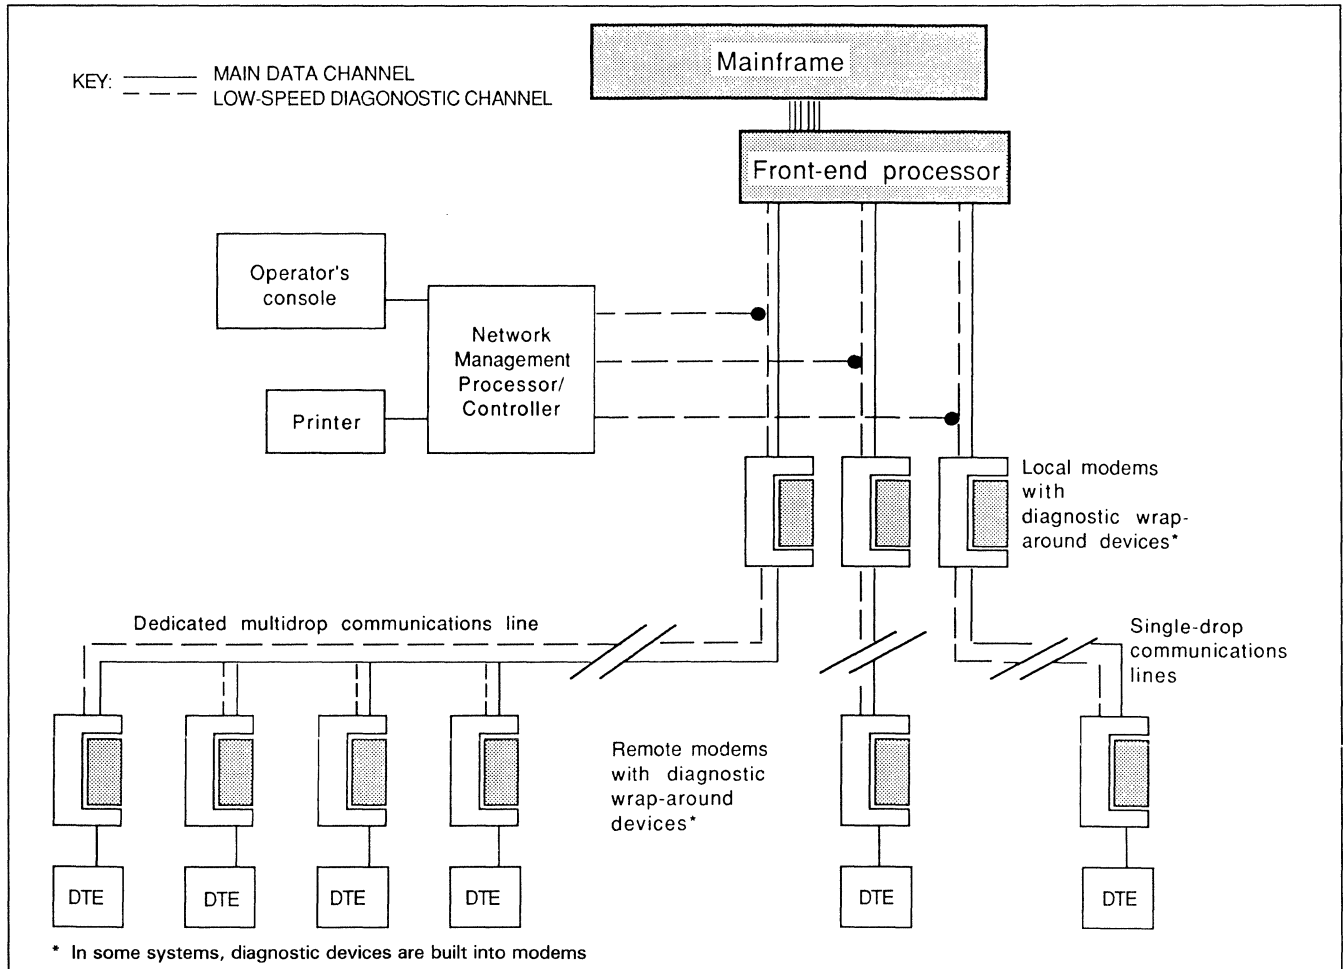


Figure 2. The typical architecture of a physical network management system.

- Accepting and acting upon error detection notifications.
- Tracing and identifying faults.
- Carrying out sequences of diagnostic tests.
- Correcting faults.

In summary, Fault Management is the discipline of detecting, diagnosing, bypassing, repairing, and reporting on network equipment and service failures.

Accounting Management—enables network managers to identify costs and establish charges for the use of communications resources. Accounting Management functions include:

- Informing users of costs incurred or resources consumed.

- Enabling network managers to set accounting limits and to associate tariff schedules with the use of resources.
- Enabling network managers to combine costs where multiple resources are used, to achieve a specified communications objective.

Configuration Management identifies, exercises control over, collects data from, and provides data to communications systems for the purpose of initiating and providing continuous reliable connectivity. Configuration Management functions include:

- Setting the parameters that control the routine operation of the communications system.
- Associating names with managed objects and sets of managed objects.

Network Management Functions: Telecommunications Hardware

- Initializing and terminating managed objects.
- Collecting information (on demand) about the current condition of the communications system.
- Obtaining announcements of significant changes in the condition of the communications system.
- Changing the network's configuration.

Configuration/name management is concerned with maintaining an accurate inventory of hardware, software, and circuits. One goal of configuration management is the ability to change that inventory in a smooth and reliable manner in response to changing service requirements. This aspect of network manage-

ment is concerned with Network Directory and related to CCITT's X.500 standard. Configuration/name management ensures consistency and validity of operating parameters, naming and addressing tables, software images, and hardware configurations of managed systems.

Performance Management enables the network manager to evaluate the behavior and effectiveness of resources and related communications activities. Functions include:

- Gathering network system statistics.
- Maintaining and examining network history logs.
- Determining system performance under normal and degraded conditions (artificially or self-induced).
- Altering network operation modes for the purpose of conducting performance management activities.

In summary, performance management is concerned with the use of network resources and their capability to meet user service level objectives.

Security Management supports the application of security policies. Security is an important issue: unauthorized or accidental access to network control functions must be eliminated or minimized. Security Management functions include:

- Creating, deleting, and controlling security services and mechanisms.
- Distributing security-relevant information.
- Reporting security-relevant events.

Security management controls access to both the network and the network management systems. In some cases, it may also protect information transported by the network from disclosure or modification. For more information on security management, see "Evolving Security Management Standards," Report NM20-500-101. ISO's document 7498-2, an addendum to the basic OSI Reference Model, also covers security issues in more detail.

LIMITATIONS OF THE OSI MANAGEMENT MODEL

These OSI definitions provide a starting point for consensus on a common nomenclature. Ideally, such a consensus may help eliminate any confusion created by the proliferation of views on what constitutes

FAULT MANAGEMENT

Network status supervision
Event notification
Alarm correlation
Problem detection
Problem determination
Diagnosis and repair
Network recovery
Trouble tracking
Network tests
Activation, deactivation, restart
Online/offline technical maintenance

CONFIGURATION MANAGEMENT

Inventory control
Configuration control
Directory management
Change management
Provisioning
Software maintenance
Service-level agreements

PERFORMANCE MANAGEMENT

Monitoring
Defining performance indicators
Reporting and statistics
Maintaining the performance database
Analysis and tuning
Procedures for fault management

SECURITY MANAGEMENT

Establishing security
Maintaining security
Partitioning

ACCOUNTING MANAGEMENT

Costing
Charging
Budgeting
Verification

Table 1. Principal network management subsystems and functions.

Network Management Functions: Telecommunications Hardware

network management and its functions. OSI categories are somewhat limited, however, and exhibit bias toward managing the hardware aspect of the network (physical devices, etc.) over software and systems management.

Datapro believes that the OSI functional categories do not adequately address the following areas.

Systems management encompasses management of not just communications links, but the applications and other software traffic that passes over those links.

Asset management assists corporations in modeling networks and determining the least expensive way to operate them.

Problem Management encompasses preparing for and handling all trouble calls received at the network's help desk. Problem management is broader than fault management, which focuses on handling alerts generated by and received from telecommunications instruments and systems. Problem management

encompasses fault management as well as network problems that are discovered by means other than alerts or events. Effective problem management includes establishing help desk procedures, audit trails, assignment groups, and other mechanisms that assist in the comprehensive handling of problems in large-scale networks.

Both vendors and users recognized the limitations of the OSI functions. For example, AT&T has claimed that an ideal network management system should provide, function in an integrated fashion, end-to-end network management including the user-allocated portion of the public network as well as private lines and customer premise equipment. In addition to the five OSI functions defined in this report, AT&T adds Planning Management, Operations Support, Programmability, and Integrated Control. (For definitions of these four added functions, see "AT&T Unified Network Management Architecture," Report NM40-313-101.) □

Network Monitoring and Control

This report will help you to:

- Implement monitoring and control measures in environments which include T1 multiplexers, T1 test equipment, CSUs, ESF facilities, and intelligent networks.
 - Prepare for the advent of ISDN in regards to monitoring and control.
-
-

According to Network Strategies, a Practice of Ernst & Whinney, 67 percent of the largest 1,750 companies now have T1/DS1 networks; by 1992 that number should grow to 80 percent. In 1989, there was a near equal split between backbone networks and simple point-to-point networks (for which network control and monitoring is relatively straightforward). By 1992, over 45 percent of the companies will have backbone multilocation T1/DS1 networks, compared with 33 percent with point-to-point networks. T3/DS3 facilities are also increasing in number.

Network management becomes much more complicated in a backbone environment; in addition to the increased port and routing considerations, multiple inconsistent network clocks create problems.

Today's networks typically include equipment from a variety of vendors, each with its own version of network control and monitoring. Neither an industry standard nor a comprehensive single-vendor solution now exists, much less a broad multivendor solution. Yet, the increasing strategic importance of information networks to corporate America has prompted managers to place more emphasis on network reliability, availability, and flexibility. This, in turn, has amplified the need for effective network control and monitoring tools and methods.¹

This report was developed exclusively for Datapro by Daniel Minoli, an adjunct professor at New York University's Information Technology Institute. Mr. Minoli is also a full-time data communications researcher and strategic planner.

T1 TESTING, MONITORING, AND CONTROL

The potential liability from either catastrophic failure of termination equipment or lines, or from a deterioration of the service due to a partial impairment, increases as more and more of the corporation's communications traffic is put through a small set of discrete facilities. The network manager must identify degraded circuits and equipment before service is totally affected. This predicament highlights the importance of monitoring and testing procedures, as well as the equipment which DS1/T1 users must implement to prevent, monitor, or resolve potential problems.

Network monitoring and control systems are typically composed of both hardware and software. The software may run on a mainframe, workstation, or dedicated personal computer and work in conjunction with network-based hardware that collects data from all network devices (such as modems, multiplexes, switches, and lines). The data is critical to the network management system in order to analyze and resolve problems and tabulate the desired performance measures into usable reports.

Many NMS workstations or consoles include a color display and offer graphical tools that present a pictorial map of the network configuration and potential problems in the time domain. Fault detection is the primary function of most network management systems; however, the ideal product provides capabilities for network configuration, performance measurement,

Network Monitoring and Control

troubleshooting, equipment control, database querying, report generation, and inventory management. Existing network management tools do not yet address all of these areas. To be most effective, a network management system must incorporate data from different vendors' components. Unfortunately, no vendor has yet developed a system that performs this entire function satisfactorily.²

T1 Test Equipment

By using test equipment to perform appropriate tests and measurements, the user can minimize the impact of degradation and failure. Test equipment isolates Channel Service Unit (CSU) problems from span line problems, analyzes end-to-end link performance, and performs preventive maintenance. Test equipment is relatively inexpensive (\$2,000 to \$7,000), except for the top-of-the-line protocol analyzers. Typically, the equipment will pay for itself (in terms of better service) in a matter of months.

Preventive maintenance of DS1 facilities can easily be accomplished with small overhead and without affecting service, given the right equipment. Test equipment can also be used from a remote location if the equipment is suitably configured.

Two types of high-capacity integrated networks exist: point-to-point networks and backbone networks.

Point-to-Point Networks: The objective of point-to-point multiplexers is to provide common route concentration. By definition, these networks involve simple topologies, even when using drop and insert. Network management systems for point-to-point networks are often manually configured and provide only basic administrator interface functionality. Network management is relatively simple and is achieved by DIP switches or rotary switches. Point-to-point networks are monitored with device-provided LEDs.

Backbone Networks: Backbone network multiplexers support user provisioning of bandwidth. The network topology of a backbone network can be fairly complex. Typically, some level of intelligence is distributed at each node to facilitate operations such as monitoring and control. High-end backbone multiplexers use PCs or workstations to collect data and/or display performance graphics.

T1 TESTING, MONITORING, AND CONTROL

Basic T1 multiplexer network management functions include circuit provisioning, fault diagnosis, alarm re-

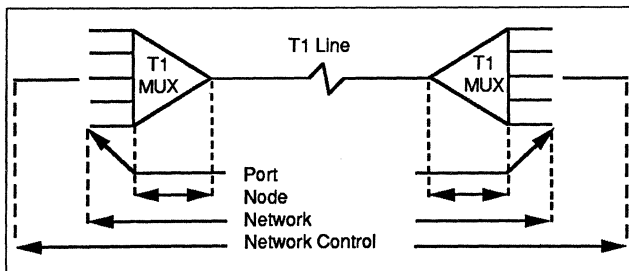


Figure 1. Control areas in T1 multiplexers.

porting, and production of management reports. Management can be done at the port, node, network, or end-to-end level, as shown in Figure 1.

Port Level. The port-level management function provides control and monitoring of the multiplexer interface to the DTE. This interface should be flexible and support test capabilities. Typical requirements include:

- Collecting performance measurements (for example, BERT, levels, lead status)
- Providing accounting reports that highlight port usage statistics
- Software control, which avoids the necessity of removing the line card and dealing with DIP switches
- Loopback testing capabilities

Node Level. The node-level management function provides control and monitoring of the node itself (the multiplexer) and of the telecommunications link. Hence, it provides the connection between the physical access port (typically an EIA-based facility) and the DS1 bandwidth. High bandwidth availability is critical; therefore, the network management system should provide a mechanism for handling dynamic bandwidth allocation. Other NMS requirements at the node include a user-friendly interface, self-test capabilities, and redundancy of key components. The node should support a local network control console and provide alarm and usage reports.

Network Level. The network-level management function provides control and monitoring at the application layer. It is responsible for application availability and quality of service.

Network-level management functions must support multiple carriers. In addition, users require easy circuit provisioning and rapid circuit restoral in case of failure. A number of high-end multiplexers provide automatic route generation; restoral time can range from 2 seconds for packet-driven multiplexers to 40 seconds or more for circuit-based multiplexers. Parameter

Network Monitoring and Control

routing, which involves sending parameters to remote locations, is also very important.

Network Control Level. The network control-level management function tracks all network components, typically provided through a relational database and a full-fledged report generator. The database must be capable of supporting complex topologies involving thousands of network components. Additionally, the database must track multiple component types (hardware, software, etc.) as well as the relationships between network components, including physical, logical, procedural, and organizational relationships.

CSU LINK MONITORING AND CONTROL

A major component in point-to-point and backbone network monitoring and control, the CSU is a protective interface designed to connect customer premises data equipment to a carrier's digital transmission line, either a Digital Data Service (DDS) facility or a DS1/T1 facility. To facilitate network monitoring, control, and fault isolation, additional built-in CSU features provide for test and monitoring access of the signals sent and received by the DTE.

CSUs provide extensive on-site diagnostics in the form of jacks, LED displays, and loopback switches. CSUs facilitate testing both the span lines and customer-premise (CPE) equipment. The latest CSU hardware allows test and monitoring functions which can be integrated with centralized Network Control Center diagnostic systems.

Service-Affecting Tests

The CSU may employ inband and out-of-band signals at the network interface to allow maintenance personnel to conduct span diagnostics locally or from a remote location. The front panel also generally provides diagnostic LEDs to display minor alarm conditions such as excessive jitter, excessive Alternate Mark Inversion (AMI) or Bipolar 8th Zero Substitution (B8ZS-coding) violations, all-ones signal detection at the network interface or CPE, and loopback status. The front panel may also provide test jacks to access, monitor, and test the CPE and the span line. These jacks are useful in facilitating test operations for quick diagnosis of communications problems. Typical signal access points on the CSU include:

- Line in, breaking facing network interface receiver
- Line out, breaking facing network interface transmitter

- Equipment in, breaking facing CPE receiver
- Equipment out, breaking facing CPE transmitter

These tests are service affecting; thus, they disrupt normal communications.

Realtime Monitoring

Other signal taps can be done while the facility is in normal operation. Jacks corresponding to these tests are:

- Line monitor, nonbreaking facing network interface receiver
- Equipment monitor, nonbreaking facing CPE equipment

Loopback tests are set locally using switch activation. Digital-code (DC) activated signals facilitate loopbacks at remote CSUs.

The ability to perform local and remote loopbacks is important, as is the ability to test and monitor the signal in both transmission directions.

Many CSUs and some T1 multiplexers can monitor T1 link performance on a realtime basis while the line carries normal traffic. Realtime tests are facilitated by implementing bipolar coding. The DS1 pulse train at output of the T1 multiplexer is encoded with a bipolar scheme before it reaches the T1 line for transmission. In bipolar coding, alternating positive and negative pulses represent one state (a binary "1"); absence of pulses represents the other state (a binary "0"). This permits the multiplexer to detect single-bit errors. If an error occurs in a 1 bit position, thereby converting it to 0, adjacent 1s will be of identical polarity, which is easily detectable, since this violates the polarity rule. If an error occurs in a zero, converting it to a 1, there will be two successive 1s of identical polarity, which also violates the polarity rule. These events are called bipolar violations. The monitoring is continuous over time and thus is very effective.

Many CSUs can also generate alarms based on selectable thresholds (for example, loss of line signal for more than 100 milliseconds). This saves time and money, as it avoids the need to connect external test equipment. Some typical parameters used include:³

- Signal level
- All-ones conditions
- Loss of synchronism

Network Monitoring and Control

- Framing error rate
- Bipolar violation rate
- Jitter
- Bit error rate
- Errored-seconds rate

Many CSUs offer signal monitoring using LED indicators, built-in loopback controls, and test signals, as well as remote control over some of these functions. Only a few CSUs, however, provide integrated centralized network control or sophisticated performance monitoring.

A new generation of intelligent CSUs may reach the market soon. These CSUs incorporate sophisticated performance monitoring capabilities, with increased diagnostic functions, and supervisory ports for communicating with the user's network control center (NCC).³ In this regard, CSUs have followed an evolution path similar to that of analog modems. Intelligent modems now incorporate monitoring and diagnostic capabilities by using a secondary auxiliary channel. The nondisruptive centralized control features in second-generation CSU equipment are similarly based on the maintenance channels provided by the ESF line format.

Some diagnostics require service-disruptive tests that insert pseudorandom bit patterns. In this context, the advantage of an advanced CSU is to provide the capability to automatically initiate loopbacks and to transmit test signals on command from the NCC. CSUs which support interfaces to the NCC may employ a dial-up approach, a polled multipoint data circuit approach, or a supervisory subchannel within the T1 link. The dial-up approach may be cost effective, but it should include security features to prevent unauthorized access. Clearly, the CSU also needs auto dial capabilities so that it can send alarms to the NCC, and auto answer capabilities to respond to NCC-initiated queries (in addition to queries, the NCC may wish to download configuration parameters).

As an example, AT&T's DATAPHONE II System Controller is a network management system that allows user to manage AT&T modems, multiplexers, and digital data sets through a single integrated system. With this system, the user can display network maps, produce trend reports, and schedule testing of network components.

The DATAPHONE II System Controller supports analog, digital, and switched network applications and two-point, multipoint, and multiplexed networks, in-

cluding T1.5 facilities. AT&T's ACCULINK Network Manager is an application software enhancement for the DATAPHONE II Level IV System Controller. It allows customers to monitor, control, test, and reconfigure backbone T1 multiplexers, an addition to the network management capabilities provided by the DATAPHONE II System Controller.¹

Recent FCC Changes on CSU Functionality

For many years, ambiguities existed concerning the type of functionality that the CSU must provide, particularly in the monitoring and control area of loopback.

On October 19, 1987, the FCC issued a News Bulletin stating it had "eliminated the requirement that carriers provide line power on 1.544 Mbps service, as well as an associated requirement that terminal equipment connected to 1.544 Mbps service contain a continuity of output capability as a registration condition under Part 68 of the rules." The deleted paragraph, 68-318(b) of the FCC Rules and Regulations, required that signals that come into the DS1/T1 network from registered Network Channel Termination Equipment (NCTE) 1) provide one of three types of keep alive signals, 2) provide pulse density maintenance, and 3) be powered by 60 milliamps from the network.⁴

The purpose of FCC Part 68 is to ensure that individuals connecting to the public network do not harm the network. Issues of concern include impact on other users, permanent damage to the network, billing protection, and access for the hearing impaired. Circuit-derived power is sometimes used to energize the critical circuitry in the NCTE; specifically, the continuity signal circuit, the pulse density, and the keep alive functions. However, circuit performance issues (e.g., bit error rate levels) are not covered in Part 68. Prior to the rule defining network connection specifications (Docket 81-216), the FCC used an interim registration specification based on AT&T's Specification 62411. This specification required such features as loopback, line powering, pulse density, and keep alive signals. Docket 81-216 adopted most of these, with the exception of the loopback requirement. (The issue was that the loopback has nothing to do with the harm-to-the-network question.)⁴

Loopback was not required by law after 81-216 was incorporated into Part 68. The Commission's original reasoning was that the market should dictate whether loopback would be required or not. Most CSUs provide loopback capabilities, since loopback is an essential network management function.

Network Monitoring and Control

In November 1985, a T1 Committee of the Exchange Carriers Standards Association (ECSA) requested that fiber optic circuits and circuits derived from pair gain systems should be exempt from the line powering requirement, because of the difficulties of providing power over fiber cables. In January 1986, the committee recommended the drafting of a standard for dry circuits (i.e., non-central office-provided power). In March 1986, five BOCs petitioned the FCC to modify the section on line powering. These changes were generally supported by the industry. On October 29, 1986, the FCC issued a news release (Notice of Proposed Rule Making [NPRM]) stating that a change, identified as Docket 86-423, was proposed, deleting all of Section 68.318(b). In October 1987, the FCC issued a report and order on the docket. In November 1987, the FCC issued an erratum identifying the specific implementation dates for Docket 86-423.⁴ With the new docket, the user must provide the power source if the service provider states that it will no longer supply the 60 milliamps after 1989. The carriers must provide power until December 1989. Any new installation occurring after February 18, 1988 may not be line powered; users may consider providing local power immediately for these circuits.

Although the BOCs only asked for the deletion of line power, in some arguments they stated that the keep alive signals were not necessary. The pulse density requirements were poorly defined in 68-318(b), and requests for clarifications had been filed. The FCC also proposed to eliminate the pulse density requirements, since some felt that this may no longer be a harm-to-the-network issue. The newer repeaters being deployed since 1987 do not have synchronization problems, although systems already in the field may indeed continue to experience the problem. In announcing its decision, the FCC said "there should be a two-year transition period for carriers to adjust to the possibility of connection of equipment without continuity of output (pulse density maintenance)." Some manufacturers could interpret this to mean that they have no requirement to maintain pulse density; however, in the existing plant, low pulse density will lead to jitter, which in turn increases the BER. Specific carriers can still tariff a transmission offering which requires that the density be maintained (as, for example, specified in PUB 62411). To achieve a viable BER, the customer equipment must continue to provide pulse density management.

Many newer NCTE models (including CSUs) already have dual-powering capabilities; thus, the power supply issue is not critical to the user. The circuit will have to be taken down for a few minutes to implement the change. Users should analyze their particular needs when selecting NCTE, from the basic Part 68 mini-

mum requirement to features such as pulse density maintenance, loopback and keep alive signals, and diagnostic capabilities.

MONITORING AND CONTROL WITH ESF FACILITIES

In the early 1980s, AT&T announced the Extended Superframe Format (ESF), a new framing format for DS1/T1 lines. ESF enhances the network manager's ability to monitor and control the backbone network. AT&T has indicated that, as Central Office equipment is replaced at the end of its useful life, new equipment supporting the ESF will be installed, and the users will be offered the new format. AT&T will continue to support the present D4 format to protect the embedded base of CPE (T1 mux) equipment.

The ESF defines 24 frames in its superframe (rather than 12 as in the traditional superframe) but only six bits in its framing pattern. This definition frees up 6,000 bps of channel without giving up any previous functionality or any additional bits. The 8,000 bps now used to manage the DS1 facility are reallocated in the ESF format as follows:

1. 2,000 bps for framing (6 bits distributed over 24 framing bits; there are 333 such 24-bit words per second).
2. 2,000 bps for error and performance determination (6 bits distributed over 24 framing bits). A six-bit Cyclical Redundancy Check (CRC) code is provided. The check character is used for monitoring transmission quality of the DS1 facility. Additional functions include false-framing protection and other performance functions implemented in the termination equipment. The CRC-6 is generated from the bits of the preceding frame (for the calculation, the framing bits of that frame are considered to be equal to "1"). This will allow the repeaters to inform the appropriate agency in the network that some component appears to be degrading. The problem may thus be fixed before total failure occurs. The CRC-6 will detect 98.4 percent of single-bit or multiple-bit errors. The CRC-6 will also provide false-frame protection, since if the CPE/CSU hardware selects the wrong synchronization boundary, the CRC will not calculate correctly. Assuming the channel was not experiencing intrinsic problems, the hardware can assume that the hardware selected the wrong boundary.
3. 4,000 bps for telemetry and facility management/reconfiguration (12 bits distributed over 24 framing bits). The data link (also called Embedded Operations Channel) becomes available for maintenance information, supervisory control, and other future needs.

Network Monitoring and Control




Figure 2. Integration of carrier-provided diagnostics.

Performance monitoring of DS1 circuits using ESF is emerging as a useful monitoring and control mechanism; it provides significant advantages compared to obtaining information via existing channel banks and CSUs. Actual digital errors can be counted, as compared to only framing errors and estimated BER with AMI's bipolar violations counts. Two different documents for performance monitoring of DS1 circuits have emerged to date: AT&T's TR54016 and Bellcore's TA-000147. The Exchange Carriers Standards Association also has issued a proposed standard (ANSI T1. 403-1989).⁵

Customers have expressed a need for information and capabilities not traditionally available to a user of carrier-based facilities. They are demanding a way to reach across the interface, and into the Central Office to access customer-specific information and controls. See Figure 2. As customers' sophistication increases, they demand these controls. The carriers are beginning to meet these needs in order to position competitively the Central Office-based services. The leading private users of T1 facilities are now starting to demand "proof of performance" from their carriers: the carriers now must demonstrate and prove that the DS1 circuit meets the performance criteria, including outage time

and BER.⁵ The features that end customers require to manage their telecommunications networks are similar to the functions that carriers have developed over the years for internal Operations, Administration, and Maintenance (OA&M).

There are several factors motivating the carriers to extend these capabilities to the users.⁶ These network management services:

- Create new revenue opportunities, while allowing reuse of already embedded OA&M systems
- Provide end users with direct, but clearly defined, control of their virtual network
- Speed up development of new services by directly extending new support capabilities to end customers
- Meet customer requirements to support a telecommunications environment characterized by many vendors, multiple equipment types, and diverse services
- Imply minimal impact on carrier operations, since service administration is centralized

Network Monitoring and Control

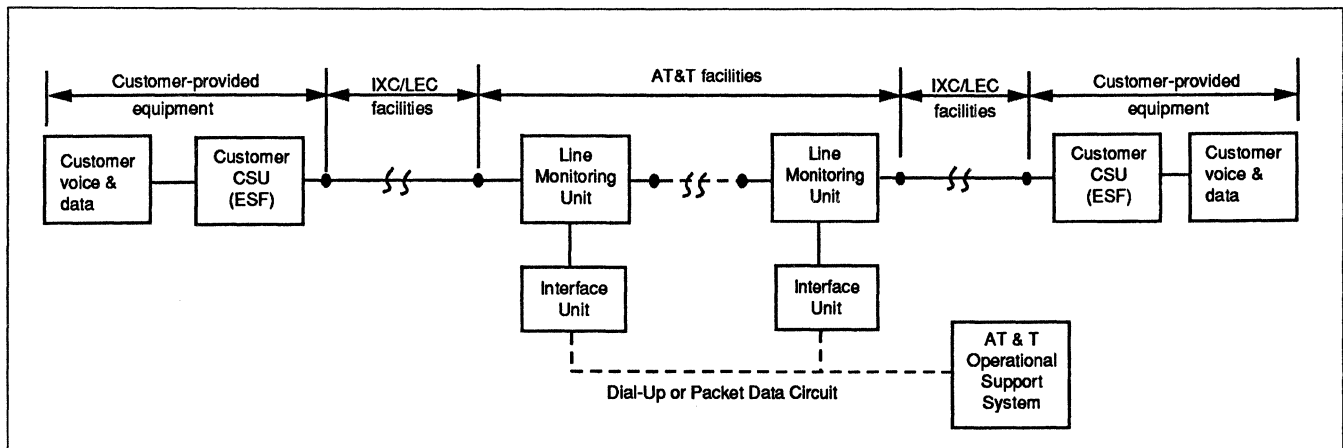


Figure 3. Deployment of Line Monitoring Units (LMUs) for T1 systems.

- Ensure consistency of OA&M information between carrier operations and end customers

In order to quickly sectionalize a problem on a DS1 circuit running through several carrier's facilities in tandem, the carrier must be capable of reading the CSU registers to check the facility between the points-of-presence. Counters, accumulators, and registers are built into the CSU to take the ESF CRC-6 and framing information and provide the following 15-minute and 24-hour performance information:

- ES (Errored Seconds): Number of second-long time intervals with one or more ESF errors
- SES (Severely Errored Seconds): Number of time intervals with more than 320 ESF errors events or out of frame (OOF) events in one second
- FS (Failed Seconds): Number of time intervals with greater than 10 consecutive SES.

AT&T has installed Line Monitoring Units (LMUs) at the point-of-presence where the T1 circuit joins its facilities. It also can add an LMU at any critical facility transfer location, to monitor the performance through its system. It can remotely loopback the CSUs and the LMUs in either direction to sectionalize a trouble, as shown in Figure 3.

While this approach is a step forward in monitoring and control, it still has limitations. The LMUs have only one set of registers, which record errored seconds in ninety-six 15-minute registers (covering 24 hours). These registers can only be reset by the carrier (not the customer) via a dedicated or dial-up facility to each LMU. The LMUs can send an alarm to the central monitoring location when errors exceed a preset threshold. The carrier and customer registers in the two terminating CSUs can also be read by the carrier via the ESF data link. The data link is queried by a

maintenance center to retrieve the information from the CSUs and the LMUs. It is also used to reset the counters, accumulators, and registers and to activate or deactivate a built-in, out-of-service data loopback testing system in the CSUs and LMUs. Messages are sent over the data link using a simplified X.25 packet proprietary (but publicly available) protocol called Telemetry Asynchronous Block Serial (TABS).⁵

Some controversy exists as to what performance information the customer should receive and whether the customer should be allowed to read the intermediate LMUs. At the moment, customers can read the registers in their own CSU and perhaps the far-end CSU if the circuit is end to end transparent to the ESF signal (it would not be transparent, for example, if the circuit went through a Digital Crossconnect). Customers cannot read the registers in the intermediate LMUs, nor can they control the loopback to sectionalize a trouble. Neither AT&T TR54016, ANSI 403-1989, nor TA-000147 envision customer access to the intermediate monitoring units.

AT&T's provision of ESF performance monitoring forced MCI to move rapidly to introduce ESF itself. MCI is implementing ESF performance monitoring on its DDS circuits, using a custom-designed TR54016-compatible system. US Sprint is also introducing ESF and ESF performance monitoring compatible with TR54016.

T1 MONITORING AND CONTROL USING DISCRETE TEST EQUIPMENT

Advanced PBXs and T1 multiplexers incorporate network management systems to varying degrees. Less functional multiplexers or channel banks require auxiliary systems and equipment. Test equipment can be categorized in the following groups: analog testing equipment, interface/access equipment, data monitors

Network Monitoring and Control

and recorders, error-rate measurements, performance measurement equipment, protocol analyzers and simulators, and centralized network management and diagnostic systems.

T1 testing can be categorized as follows:

	Analog	Digital
Data-oriented	x	x
Voice-oriented	x	x

Voice-oriented analog testing generally relates to levels, bipolar violations, and so on. Voice-oriented digital testing generally involves clocking/synchronism, signaling bits, VF drop/insert, errored-seconds, and so on. The basic need of most voice testing personnel is the ability to access one or more channels without interrupting service; other needs include monitoring for loss of frame or loss of loop and the ability to signal or simulate signaling functions.

For years, telcos have performed this type of testing to manage the plant. Data-oriented testing follows the same principles, except that the drop/insert of 64K bps subchannels is not a major concern; the aggregate nature of the DS1 signal and the ones-density issues are more prevalent.

In-service monitoring can be used to test facilities periodically so that degraded conditions can be identified before they affect service. When traffic-carrying circuits cannot be disturbed, in-service monitoring is desirable. In some cases, the traffic could be rerouted to perform out-of-service testing, but the cost involved in this action (either in transmission or labor charges) may be prohibitive. In-service monitoring can also precede out-of-service testing, since localizing potential problems with passive monitoring may reduce the out-of-service repair interval.

In-service testing can be divided in two classes: 1) those measurements that apply to the DS1 signal proper; 2) measurements applying to the content of the information carried on the DS1 facility. The former includes coding violation analysis, excess zeros detection, signal frequency determination, signal level, all-ones condition, and time jitter analysis. The latter includes framing error analysis, cyclic redundancy check analysis, and yellow alarm detection.⁷

Table 1 depicts typical functions of a T1 tester.

Line errors are introduced by externally generated noise (motors, electrostatic discharges, and so on), mechanical failures of the cables or repeater equipment,

improper engineering (loss plan, etc.), faulty user equipment such as CSUs and multiplexers, and so on. In general, there are nine different types of analog line impairments that affect the performance of a link:⁸

Noise—inherent to the line design or induced by transient energy bursts. There are several types of noise: 1) Gaussian noise (also called white, thermal, or shot noise); 2) impulse noise hits; 3) cross talk noise; 4) ground-loop or common-mode noise; and 5) quantizing noise. Some of these noises impact a given transmission medium, and some impact other transmission media.

Attenuation—the loss of signal level as it passes through the transmission line.

Attenuation distortion—signal-level loss that affects the relative magnitudes of various frequency components of the transmitted signal.

Envelope delay distortion—phase nonlinearity where the delay is greater at the edges of the passband, caused by inductive and capacitive reactance in the copper network. Such problems can occur in a copper line with coils.

Frequency translation—the shifting of all frequency components in the modulated signal.

Jitter—of three types: phase jitter, amplitude jitter, and time jitter. Phase jitter results in a pure tone having an associated FM spectrum. It is caused by coupling from power line-associated equipment and ringing generators. Amplitude jitter results in a pure tone having amplitude modulation. Time jitter is caused by digital repeaters and clock drifts.

Intermodulation distortion—sometimes called nonlinear distortion, the difference between harmonic distortion and IMD is the method of measurement.

T1 TEST EQUIPMENT FUNCTIONS	
Analyze framed or unframed DS1 data	
Measure errors, BER, and error-free seconds	
Measure bipolar violations, violation rate, and violation-free seconds	
Generate up- and down-loop codes for CSU control	
Connect to the DS1 span line through a CSU or at a cross-connect point	
Use source timing or can derive timing from loop timing	
Allow testing through a DCS	

Table 1. T1 testers provide a variety of functions. By using test equipment to perform appropriate tests and measurements, the user can minimize the impact of network degradation and failure.

Network Monitoring and Control

Single frequency interference— the addition of one or more discrete frequencies to the signal as it passes through the line. For example, a copper line near AC power lines would result in 60 Hz basic and harmonics being added to the signal.

Transient—impairments that do not continuously affect the line. They come and go. Dropouts are transient impairments that result in sudden reductions in the signal levels that last for more than several milliseconds; gain hits are sudden increases in signal level that last for several milliseconds.

While these types of problems are more prevalent with copper and microwave systems, all digital systems are affected by time jitter, clocking problems, and attenuation. (Analog tests are still necessary in systems with digital modulation, as long as the underlying medium is analog.)

In-Service Testing Equipment

Access Testers. These testers provide VF output with voice signaling bit status for any 1 of the 24 channels of the D4 DS1 bit stream. The VF signal can then be further tested with other analog equipment. This equipment may also give a "missing carrier" indication and "frame misalignment" indication. Access testers are ideal for voice applications but can also be used for data applications in those cases where the testers provide additional functions such as jitter measurements, BERT, and so on. Add/drop-type tests are greatly simplified with the use of access testers.

There is a growing need for channel access test equipment for integrated voice and data. In addition to providing access to the DS0 slots and to the A/B signaling bits, this equipment may also provide monitoring for error performance and integrity checking and bit pattern synthesis. Some of the equipment will also operate on DS1C signals (two DS1 streams). A number of models also offer local and remote loopback capabilities. Remote loopback allows a remote test unit to activate the loopback via loop-up and loop-down digital codes embedded in the DS1 stream. (Loopback can also be achieved in other link components, such as the CSU, so that a piece of test equipment is not mandatory to obtain this functionality.)

Only a few of the available access testers offer ESF testing capabilities, but it is likely more vendors will provide this capability within the next two years. ESF equipment provides VF output with A, B, C, and D signaling bit status for any 1 of the 24 DS0 subchannels in an ESF DS1 facility.

Some of the available access test equipment offers software configurability, allowing control of the test equipment from a test console (or from a terminal incorporated directly in the tester). These testers are generally menu driven, and test specifications and measurements can be preprogrammed in software. The base price of access testers is in the \$2,500 to \$7,500 range; optional features can raise the price into the \$4,000 to \$15,000 range. Increased complexity of the equipment, and the software programmability from a PC, will drive up the price of this equipment over the next few years.

Bipolar Violation Testers. As indicated above, noise and cross talk can affect the bipolar signal to the point where a no-pulse condition is interpreted as a pulse condition. Bipolar violation tests can count these incidents. A high bipolar violation rate generally indicates an intrinsic problem. An instrument that measures this data can be useful in sectionalizing the potential problem. Bipolar violation measurements are ideal for interLATA and interLATA T1 facilities (and less useful for nontelco facilities). (NOTE: The B8ZS method employed to achieve clear-channel conditions uses deliberate violations. Users must appropriately configure the equipment testing B8ZS lines, so that these intentional violations are not counted as errors.)

Excess Zeros/Framing Errors Testers. Framing errors occur when the receiving equipment (or a repeater) is unable to recover the clock. This condition can be caused by a number of system impairments. Many older repeaters in DS1 lines require a specific density of one-bits to maintain synchronism. While the CSU should guarantee that the appropriate rate is maintained, not all CSUs explicitly provide this function. Excess zeros can cause the repeaters to shut down and put the facility out of service. Hence, excess zeros detection instrumentation can help resolve or prevent problems. Note that the new FCC rules no longer mandate the ones-density requirement; however, the ones-density ratio should still be maintained for performance reasons. Loss of synchronism leads to framing errors, where an incorrect bit appears in the framing position (193rd bit); the incorrect framing word of 12 bits precludes the subchannels to be correctly decoded. These testers accumulate and display DS1 frame loss and framing bit errors; they may also have separate digital counters to allow the user to accumulate frame loss and framing bit errors separately. Frame loss rules are such that two out of four successive framing bits in error determine a frame loss. One aspect of the framing error measurement is that is approximates the actual BER of the facility, even for high error rates. Thus, it can serve as an in-service approximation of the end-to-end performance of the facility without having to do the out-of-service testing. Tests

Network Monitoring and Control

that monitor errored seconds provide information of the distribution of errors over the test period.

A reasonably clean T1 circuit should be 95 percent error free over a 24-hour testing period (the total number of testing seconds minus the errored seconds that occurred). The typical price for DS1 frame analyzers ranges from \$3,000 to \$6,000, depending on sophistication.

Signal Levels Testers. The DS1 signal must satisfy certain defined electrical levels at the repeater, CSU, and DCS points. Signals exceeding the spec will generate cross talk on other circuits; low signals will be masked by noise. It is important to be able to measure analog signal levels.

Jitter. Jitter is a displacement in time of the signal transitions compared to the ideal signal. Whether provided by the network or by the user equipment, timing must be within 50 parts per million (80 Hz at DS1 rates). Operation outside this range will result in jitter, which in turn implies timing slips and data errors. It is important to be able to measure jitter. The severity of the distortion depends on both the amplitude and the frequency of these displacements. The predominant source of jitter is multiplexers and regenerative repeaters.

Network and user equipment will have a certain tolerance to jitter; this will depend basically on the sophistication, thus cost, of the timing recovery circuits. How well the jitter is controlled (and how much the various DS1 components can tolerate the jitter) is important to achieve a robust and reliable backbone T1 network. The spectrum analyzer equipment can provide a detailed picture of the jitter problem. However, this type of equipment is expensive: the combination of the jitter demodulator and the spectrum analyzer can range from \$25,000 to \$35,000. Additionally, the equipment is not truly portable. Only the most sophisticated installations use this type of equipment.

Yellow Alarm. This indication is transmitted to a device when another piece of equipment detects a problem in the signal received from the first device. For example, a Digital Crossconnect System (DCS) may send a yellow alarm to a user's T1 multiplexer, if the DCS detects some anomaly. It is useful to detect yellow alarms, as these are directly indicative of a potential problem. Many testers will detect this condition.

All-Ones Condition. All-ones signal patterns are transmitted to keep the repeaters synchronized, even when no real data is being transmitted. Detection of this alarm condition when it should not exist will indicate a failure of some network component.

Cyclic Redundancy Checking. The ESF format allows users to conduct more sophisticated in-service tests. Some of these newer tests are based on a CRC code embedded in the bit stream formed by the accumulation of the 193rd bit. The CRC provides a measurement of line quality.

Service-Disruptive Testing

Sometimes out-of-service testing is the only recourse available. These types of tests may involve two pieces of test equipment, one at each end, or a single instrument connected through a loop-around arrangement.

Bit/Block Error Rate Testers. These perform the functions of a traditional BER/BLER tester but operate at the DS1 rate. Pseudorandom bit patterns are transmitted down the line and recaptured by the tester on the receive side of the looped-around line, for comparison and error assessment. Direct BER tests must be done while out of service, particularly when sectionalization of a problem is required.

An approximate method involves accessing the CRC data provided by ESF, if this service is available in locations where the user is based (actually available at both ends). The second method involves rigorous out-of-service BER testing. Performance rules of thumb are as follows:⁹

- A BER of 10^{-5} is acceptable for voice communication; a BER of 10^{-4} is marginal, and 10^{-3} is unacceptable
- Video requires a BER of 10^{-6} or better, depending on the encoding technique employed
- Data transmission requires at least 10^{-6} at relatively low data rates and 10^{-7} or better at higher data rates

In addition to bit errors, a full-featured BER tester can also detect and generate bipolar violations, as well as determine the error-free seconds (over some horizon, typically a couple of hours).

ISDN MONITORING AND CONTROL

A number of *Fortune* 500 companies are already starting to install ISDN lines, and ISDN test equipment and other products are now reaching the market. A number of such products were displayed at COMNET '89 in Washington, DC.

Network management capabilities are critical for effective ISDN implementation. There are three levels of management that users need and which ISDN must

Network Monitoring and Control

provide: receiving alarms; controlling, altering, and reconfiguring private network resources; and using private network management systems to control the configuration of publicly provided ISDN services.¹⁰

International standards organizations did not start to focus on ISDN network management until mid-1987. The CCITT Red Book, published in 1984, barely mentioned network management. In contrast, network management will be a focus of ISDN recommendations in the Blue Book, to be issued later in 1989.

Network management will be the single important area of study for CCITT in the next four years. One of the important questions now on the standards bodies' agenda is whether the ISDN will adopt the OSI Common Management Information Protocol (CMIP) or adopt something under the ISDN message set (Q.931) umbrella.¹⁰ It is generally accepted that the D-channel will be the mechanism for the management of narrow-band ISDN lines.

The North American ISDN User Forum (NIU) is now looking into the following areas: 1) ISDN Network Management user requirements; 2) existing ISDN and pre-ISDN network management systems; and 3) emerging ISDN standards. Work is expected to continue for the foreseeable future. In the past, management issues have not been addressed until the end of the standards development process. These efforts attempt to reverse the process.

ISDN network management is divided into two topological areas: managing CPE and managing the network. Common Channel Signaling System 7 (SS7) is regarded as the vehicle for managing the network: SS7 provided node-to-node control, and it will enable ISDN switches to communicate with each other. The network can also assist the CPE in its customer-oriented network management task. The 2B1Q line coding specification used in ISDN contains a Cyclical Redundancy Check and a Far End Bit Error (FEBE) check.¹⁰ This allows the network to offer testing and diagnostic capabilities to the user, similar to those available under ESF.

There are between 50 and 100 features a user might want to control and monitor with analog Centrex. With ISDN, that number could double. The ISDN user will be dealing with channels that can be allocated dynamically. Both AT&T and Northern Telecom are in the process of deploying ISDN network management capabilities into Centrex. AT&T named its product NetPartner; Northern named its system Business Network Management (BNM). These products are now in development and are being tested at customer sites. Commercial availability is expected by the summer of 1989.

CARRIER-BASED CONTROL SYSTEMS WITH INTELLIGENT NETWORKS

Carrier-based intelligent networks first appeared to enable carriers to introduce new services quickly without having to modify the hardware and software of a large number of switches. One example of an intelligent function is the 800 Service, where a logical number is translated into a physical address. The proliferation of intelligent networks is continuing and will accelerate in the early 1990s. By centralizing the service control, intelligent networks provide for the ability to reconfigure the network in realtime.

Intelligent network capabilities can be applied to provide sophisticated network control and monitoring functions. Typical performance requirements of a carrier-based network management and control system for end-user access are as follows:¹¹

Percent of undetected user faults—the faults detected by the user rather than by the system. No more than 5 percent of the total faults should be detected first by the user rather than by the system.

Response time to fault isolation—the time it takes to isolate a fault once it is reported. The system should have the capability to isolate faults to the source or cause in typically five minutes or less.

Response time to resource reallocation—the time it takes for the network to respond with backup resources to allow for continued operation, which may vary from user to user. Some users may be willing to pay for total redundancy in the network, and its time may be a fraction of a second. On the other hand, users may not have direct redundancy in the network, and alternate routing may be necessary.

Response time to user complaint—the system must be capable of responding to a user query in less than several seconds. This includes only the answering of the phone for a typical request. In addition, the system must provide the customer service representative with the direct access to the end user and of all information on his or her account.

In addition to the basic intelligent network technology, emerging international standards will facilitate user access to carrier-provided diagnostics. Of particular note are the studies being conducted on the Telecommunications Management Network (TMN) architecture by CCITT in Study Groups IV and XV.

Network Monitoring and Control

Telecommunications Management Network (TMN) Architecture

According to TMN principles, network operations and management strategies, in addition to end-user network management, depend on three key elements of network technology: 1) standard operations and management functions in network-deployed Network Elements (NEs); 2) discipline-independent network support applications, which implement standard functions and interfaces in an open architecture; and 3) industry-standard interfaces between applications and NEs.

The purpose of a TMN is to support a carrier in the operations and management of its telecommunications network. The basic concept is to provide an organized structure to achieve the interconnection of various types of Operations Support Systems to telecommunications equipment (Network Elements) using an agreed-upon architecture with standardized interface configurations and protocols. While principally for carriers, the NE data can prove useful for monitoring and control and can be made available to the user, employing a standardized interface. The TMN is not part of the public telecommunications network but is functionally a separate network that overlays it and interfaces with it at a number of points to receive information from it and to control its management operations.¹² The TMN architecture will provide carriers and equipment manufacturers a set of standards to use in designing a management network and in developing equipment for a modern telecommunications network.

TMN principles were first expounded in CCITT Draft Recommendation M.30 (M.2x). The TMN provides the means to transport and process information related to the management of a telecommunications network. The TMN comprises: 1) Operations System Functions (OSFs), (2) Mediation Functions (MFs), and (3) Data Communications Functions (DCFs), which include both local access and backbone communications. The TMN is also connected to Network Element Functions (NEFs) and Workstation Functions (WSFs). These functional components are independent of possible physical configurations. Standard interfaces defined correspond to conceptual points of information exchange and include:¹²

“Q” Interfaces between OSs, MDs, and NEs. The Q Interface has three levels, providing flexibility in design and implementation:

- “Q1” connects NEs without MFs to MDs

- “Q2” connects devices (NEs or MDs) which contain different levels of MFs via the Local Communications Network (LCN)

- “Q3” connects MDs, NEs with MFs, and OSs to OSs via the Data Communications Network (DCN)

“F” Interfaces between TMN components and WSs.

“G” Interfaces between the WS and the user.

“X” Interfaces between the TMN and other networks/TMNs.

Standard interfaces for the connection of real open systems, such as applications and NEs, are being defined in terms of the Open Systems Interconnection (OSI) Reference Model. Application messages are also considered as application protocols. These application messages/protocols provide the means for an application process involved in a distributed information processing task to communicate information to a peer application process to complete the information processing task. Currently, the “Q” series of interfaces are a primary focus among the standards bodies. The TMN is intended to support a wide variety of functions that are used to perform the operations, administration, maintenance, and provisioning (OAM&P) activities of a telecommunications network. Some of the most important functions are:

- Performance Management
- Fault Management
- Configuration Management
- Accounting Management
- Security Management
- User-System Interface

LOCAL AREA NETWORK MONITORING AND CONTROL

Managers, operators, and diagnostic applications require timely and accurate information to optimize LAN performance. While it is generally agreed the LAN management involves monitoring, control, and diagnosis, there is no universally accepted definition for LAN management.¹³

Historically, LAN system designers did not recognize management as a major requirement, since the need to share resources is the single major driving force behind LAN connectivity. For commercial reasons, manufacturers have promoted connectivity, accelerating the growth of complex systems without regard for their manageability.

Management strategies are now constrained by the limitations of the installed technology. The same

Network Monitoring and Control

forces that drive standalone users to connect with LANs also drive LAN users to connect with another heterogeneous network; any LAN management strategy must consider complex environments into which all LANs are evolving. The historical approach to LAN management is to identify a need, then develop a utility on a case-by-case basis. It is likely that many commercial utilities were originally conceived by test engineers to help them in their duties and were then marketed to end users.¹³ The problem is not in the utilities themselves, but in the approach. Management support should be built into LAN architectures, rather than derived from a collection of utilities.

The current arsenal of utilities for LAN management includes line monitors, logic analyzers, management applications, databases, and expert systems. Each of these tools fills a need. For this reason, system designers must build frameworks that allow customization. In LAN management, information must be available from the lowest layers of communication; hence the use of line monitors and protocol analyzers. LAN management facilities must track an array of data, such as network topology, security levels, and accounting; this mandates the use of databases. LAN management is an expert-intensive function; with a shortage of experienced personnel, expert systems (discussed below) may play an important future role in managing complex LANs.

Several factors drive the monitor and control process:

Information must be gathered continuously.

Information must be gathered from all network components, at all layers. The monitor and control facility must be distributed throughout all OSI Layers of the LAN, as already demonstrated by the OSI Network Management framework. Through this set of Layer Management Entities, the management system acquires information concerning LAN status. It may also exercise control over LAN components such as initiating diagnostic procedures. A diagnostic Application Program Interface (API) from Layer 7 to the application that will analyze the data must be provided by the LAN communication resources. This diagnostic API provides a window through which monitoring applications may extract information such as driver status, routing tables, and file server connection status. Making such information available is obviously the foundation of any effective LAN management system.

The monitor and control facility must use network communications resources, since it is not practical to have a separate management network for the LAN. A separate network would add to the cost of the LAN and would also have to be managed. Since the monitor and control facility uses the network, it consumes network re-

sources; the adverse impact of this traffic on LAN performance must be minimized.¹³

Monitoring the LAN is the most demanding function of the management facility. Two methods are available for monitoring a LAN: polled and spontaneous. (For more information on polled vs spontaneous monitoring, see "Managing Local Area Networks: Accounting, Performance, and Security Management," Report NM50-300-501.)

In a LAN environment, data can be classified as either static or dynamic. Static data pertain to LAN configuration (for instance, the type of file server or workstation). Dynamic data pertain to LAN operation (for instance, dynamic routing tables and congestion statistics). A sophisticated monitoring and control facility can provide operators with all needed static and dynamic data. The function of the database facility is to store this data. Several constraints affect the database facility; in particular, management databases must deal with congestion, data reduction, and aging. The stream of information destined for storage in the database facility consumes communication bandwidth. The database facility may be distributed to disperse management traffic throughout the LAN. The data can also saturate LAN storage capacity. This requires a strategy for reducing the amount of information stored in the database. For instance, storing minima, maxima, averages, and current values.¹³

MONITORING AND CONTROL WITH EXPERT SYSTEMS

Knowledge-based expert systems are finding applications in network management and control. Some of the areas where knowledge-based expert systems apply are diagnosis (fault management), fault isolation (fault management), preventive maintenance (fault management), network design (configuration management), network operations (combination of all five functional areas), routing, user-friendly interfaces, communications software development, and long-range planning.

Despite the availability of advanced AI tools, the technology is by no means off the shelf. Developing useful and cost-effective applications remains a challenging and time-consuming task. Thus far, expert systems have had mixed results in actual field trials. In the most successful cases, the project was kept simple and the goal was to eliminate the tediousness of mundane tasks, thereby reducing management labor costs. The most promising (and also most ambitious) long-term potential for expert systems is in analyzing the interrelationship of network events, particularly in the role of cause-and-effect diagnostic analysis.¹⁴

Network Monitoring and Control

Fault Diagnosis and Maintenance. Fault diagnosis and maintenance were the first two areas in computer networks where AI techniques found applications. Expert systems can handle large volumes of complicated data. If successful, these systems can increase network reliability (uptime, path quality, and so on) and provide the ability to diagnose and resolve problems more expeditiously.

Some of the challenges encountered in this area include the lack of international standards and the high sensitivity of the system to message formats. Expert diagnosticians analyze network failures and impairments to determine the probable causes and then specify the required repair and maintenance to resolve the problem.

Network management diagnostic systems monitor irregular behavior in an effort to arrive at its cause. Such a system analyzes network outage data to determine the cause and then initiates appropriate dispatch and repair (human, or even automated). Preventive maintenance expert systems (sometimes called control expert systems) interpret the current state of the elements under their jurisdiction and attempt to predict future states. This type of system is an example of proactive network surveillance. The expert systems review network events as they occur. They also clarify these events to a human operator by assigning a specific meaning that is understood more easily by the human than some cryptic message.

The idea of "expert operators" might be a valuable tool in data communications networks. A number of prototype AI diagnostic systems for switch maintenance have appeared in the past few years. An example of a diagnostic system in the telecommunications switch environment is Real-time Expert Analysis and Control (REACT)/Switching Maintenance and Analysis Repair Tool (SMART) system. Other diagnostics systems have been developed for LANs and DDS.¹⁵

Network Control. Large networks require substantial resources to run effectively, including instrumentation and human operators. The goal of expert systems in this arena is to assist the operators, not to replace them. The benefits/goals of a network control expert system encompass the following: increasing the accuracy and efficacy of the operator intervention; reducing the amount of ancillary information required by the operator; facilitating the operator's decision-making process; and reducing the amount of time required to restore or alter the network.¹⁵

Network Interpretation. The challenge of event interpretation in the existing environment is that operator support programs only provide status messages without any internal interpretation. Even the priority cod-

ing associated with some of these reporting systems fails to convey the true state of affairs—operators interpret a group of messages rather than interpreting individual messages. Individual messages arriving at a console have priority orderings. The expert system must take the ordering into account to arrive at the correct conclusion. Another area where expert systems can help the network operator is in providing graphical displays of the network resources.

Network Design. Engineers have advocated AI systems for computer-aided modeling and performance evaluation of communications networks for a number of years. Data communications networks require frequent reconfiguration. Network reconfiguration tools associated with some of the more advanced network management systems have alleviated this problem to an extent; however, the promise of expert systems is to make this design effort even more expeditious. A number of prototype expert systems for network configuration are available. These systems are found in two forms: 1) a consulting service or 2) an end-user application.¹⁵

REFERENCES

- ¹J.B. Brinsfield, W.E. Gilbert, "Unified Network Management Architecture—Putting it all Together," *AT&T Technology*, Volume 3, Number 2, 1988, pages 6-17.
- ²D.M. Weston, *Network Management: Issues, Products, and Challenges*, SRI International Report D89-1318, February 1989.
- ³M. Lefkowitz, "Intelligent CSUs for Performance Monitoring," *Telecommunications*, November 1986.
- ⁴W.J. Buckley, "T1 Standards Battles," *TE&M*, January 1988.
- ⁵E.V. Hird, "A Decision is Needed on T1 Performance," *TE&M*, January 1, 1989, pages 39-45.
- ⁶G. Hickman, K.A. Perry, "An Integrated Network Management Prototype for End Customer Control," *IEEE 1988 Network Operations and Management Symposium, Proceedings*, IEEE #88CH2532-0.
- ⁷R.C. Troutman, "T1 Circuits: In-Service Monitoring and Out-of-Service Testing," *Telecommunications*, July 1986.
- ⁸J. Douglass, "How to Find Phone-Line Faults and What to do about Them," *Data Communications*, September 1988, pages 179-197.
- ⁹B.W. Worne, "An HMO for T1 Circuits," *TE&M*, November 15, 1986.
- ¹⁰F. Knight and N. Morley, "Will ISDN Address Network Management Issues?" *Business Communications Review*, March 1989, pages 78-83.
- ¹¹T.P. McGarty, L.L. Ball, "Network Management and Control Systems," *IEEE 1988 Network Operations and Management Symposium, Proceedings*, IEEE #88CH2532-0.
- ¹²T.M. Bauman, R.C. Boyd, "OS/NE Interface Motivation and Objectives," *IEEE 1988 Network Operations and Management Symposium, Proceedings*, IEEE #88CH2532-0.
- ¹³D.S. Wolston, D. Neibaur, "Architectural Support of Network Management: An Alternate View," *IEEE 1988 Network Operations and Management Symposium, Proceedings*, IEEE #88CH2532-0.
- ¹⁴M. Pyykkonen, "Network Management: End-User Perspectives," *Telecommunications*, February 1989, pages 23-24.
- ¹⁵E. Ericson, L. Ericson, D. Minoli, *Expert Systems Applications to Integrated Network Management*, Artech House, 1989. □

An Interactive Network Display System for Network Management

This report will help you to:

- Define the objectives of an interactive network display system.
 - Represent network structures graphically in software.
 - Employ the ideal network operations architecture that supports network display systems.
 - Review examples of graphically oriented network display systems.
-

The display of network operational data in graphical form has been a major challenge for network management systems. As networks are becoming more complex and as the amount of operational information increases, there is a growing need for network management systems to provide functional network display capabilities which have the ability to filter and present to network operators, in near real time, data gathered from a number of different network elements.

This report examines the objectives of an operationally functional network display system, describes the approach taken to represent a variety of network types—the basis for a specific network display system (NDS); discusses the network architecture required to support network wide data collection and display and gives examples of NDS capabilities. The report concludes by comparing the network display system objectives with the NDS implementation.

NETWORK DISPLAY SYSTEM—OBJECTIVES

- Display of Many Networks

This Datapro report is based on “An Interactive Network Display System for Network Management” by J.M. Meunier, Bell-Northern Research, Canada. © IEEE 1988. Reprinted by permission from *IEEE 1988 Network Operations and Management Symposium*, New Orleans, LA, February 28-March 2, 1988, pp. 1.18-18.18.

- Display of Large Networks
- Near Real Time Updates of Network Status
- Interactive Query of Network Information
- Interactive Control of Network Elements

Since telecommunications networks are, in reality, collections of varying numbers of cooperating networks, a basic objective of a network display system is to be able to depict, using a consistent methodology, a number of different but related networks. For example, corporate networks usually consist of voice and data networks, leased and nonleased facilities, etc. Therefore, a network display system must have the ability to display different networks under a single network manager’s control.

A network display system must also have the ability to display effectively both small and large networks (10 to 5,000 nodes).

Index to This Report	Page
Network Representation Technology	202
Network Display System Architecture	205
Conclusion	210

An Interactive Network Display System for Network Management

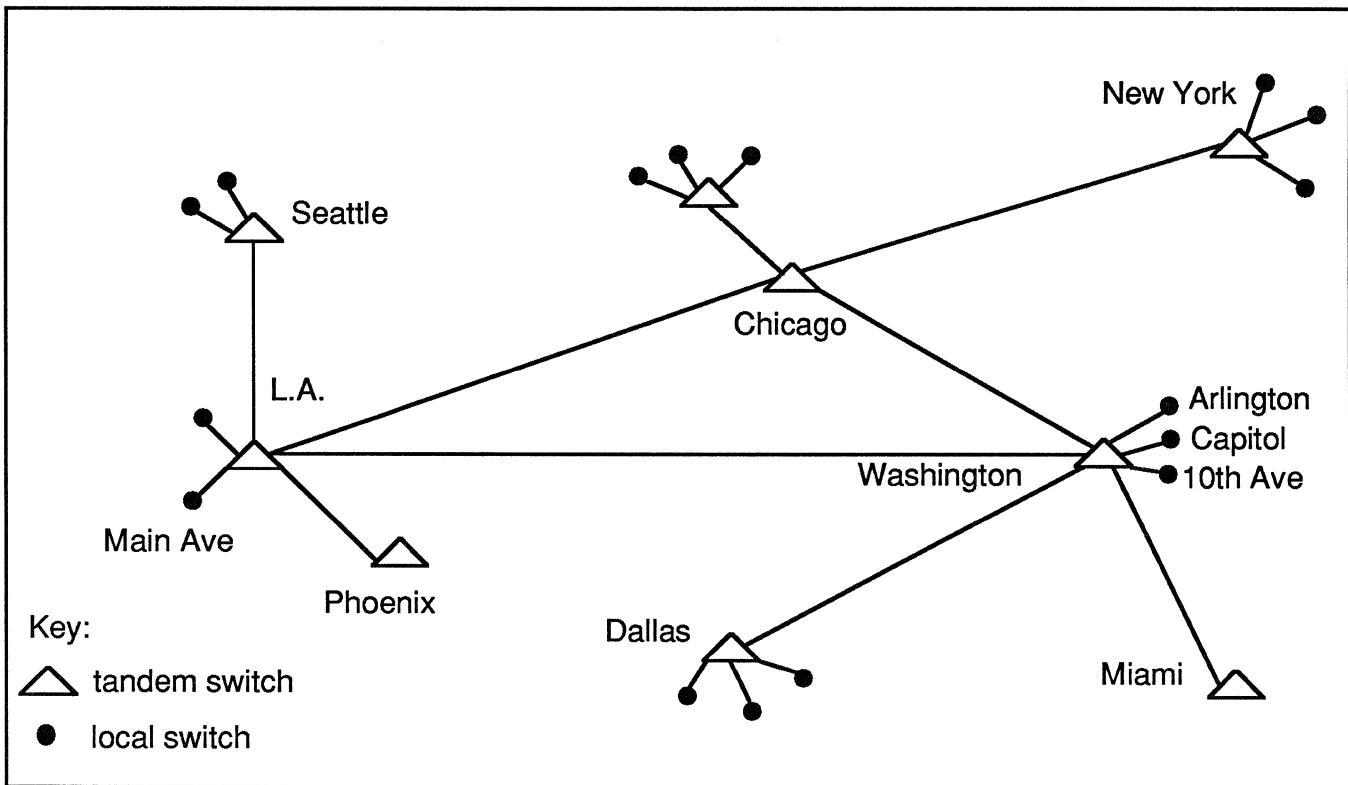


Figure 1. Unique network representation.

The network displays will be useful only if they can depict up-to-date views of network status. Techniques to clearly distinguish between different types of traffic and equipment status conditions are required.

Although the above objectives of displaying the status of a number of large networks meets the strict requirements of network display, such a system would not be operationally functional without providing the network operators with the capabilities to query interactively the network for more information—more detail about a specific trouble spot or historical information. Moreover, the network display systems must provide a window into the control of network functions such as routing, service orders and so on.

Since it is clearly impossible to show all network links and nodes for large networks on one panel or screen at the same time, the challenge of network display is to logically partition network information in a way which is appropriate for display purposes and which is meaningful to network managers. Therefore, a key technology is the representation, in software, of network structures. Three distinct approaches are described in succession.

NETWORK REPRESENTATION TECHNOLOGY

The first approach, which we call unique network representation, is a common method adopted to display networks and consists in the logical relationship between different classes of nodes—in other words, the homing hierarchy.

The usefulness of this method is limited as it is a simplified abstraction of the network being managed—the nodes and links in the diagram have ambiguous meaning (see Figure 1). For example, it is not clear that the lines are representing trunk groups, data circuits, transmission systems or physical routes. A variation of this approach is to substitute the homing hierarchy by the physical connectivity which often reflects the homing pattern.

A second approach to represent networks uses the spatial relationship which exists between different networks at the node level (a network can be a node in another network). This method satisfies one of our objectives, namely managing the size of networks being displayed. By judicious partitioning, using geographical, political and administrative boundaries very large networks can be represented by successively nesting networks within a higher level abstraction of a node.

An Interactive Network Display System for Network Management

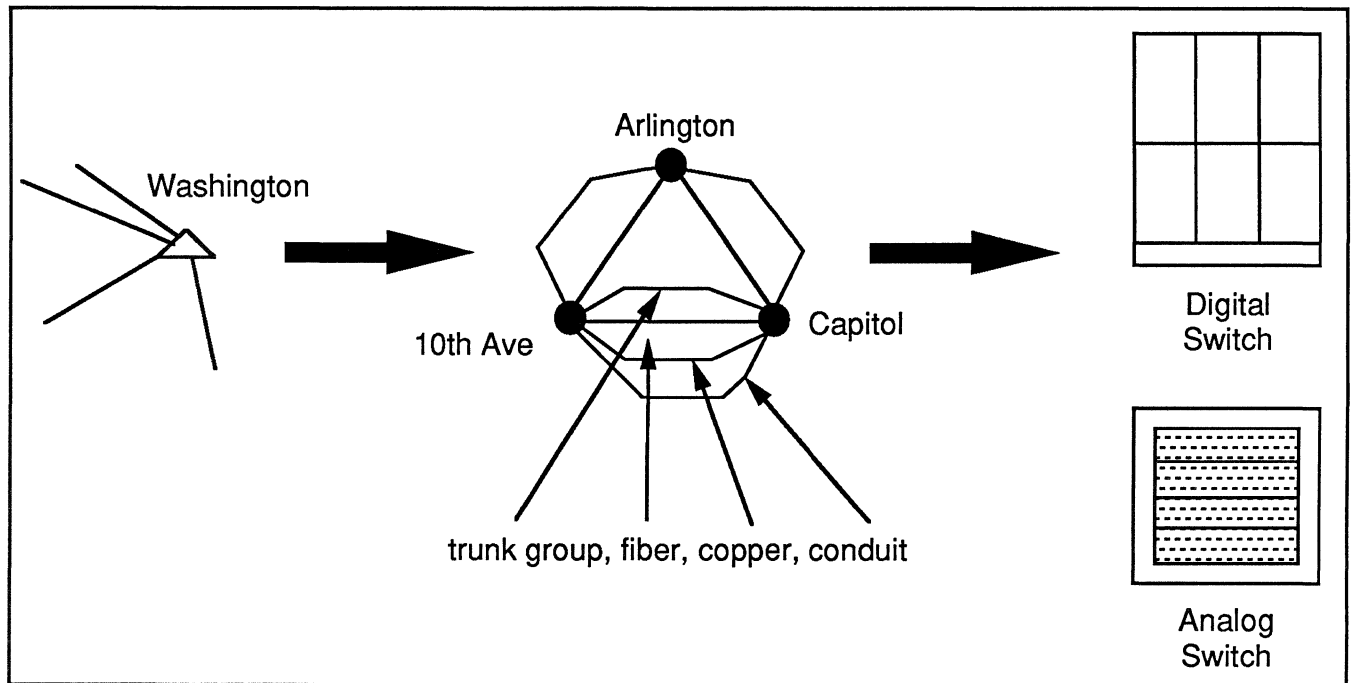


Figure 2. Node based network nesting.

In our example (see Figure 2) the nationwide Washington node consists of three switching centers; in turn, one of the switching centers is made up of two switches. Therefore, three different networks would be displayable: country-wide, city-wide and the network which exists within the switching center.

Typically, implementations of this scheme do not include link nesting—i.e., all links between nodes are shown. For example, there are four links between switching centers 10th Ave. and Capitol, a trunk group, a fiber transmission system, a copper transmission system and a conduit; each link has its own characteristic, individual status and associated operational procedure. Although this second approach meets the objective of displaying large networks, it would not be well suited to a large number of link types in a highly connected network.

Conceptually, networks consist of nodes and links—nodes are objects which are termination points and links are objects which terminate on nodes and interconnect nodes to form a network.

A third approach to network representation is based on the principle that nodes and links which share common characteristics or functions can be logically grouped to form a network layer. The association of all network layers forms the overall network.

For example, a telephone set is connected to another telephone set through a voice call link (see Figure 3). In the call layer of the network, telephone sets are

nodes and calls are links. The call layer depends on the trunk layer of the network for logical connectivity. The trunk layer consists of switches and trunk groups. In turn, the trunk group layer relies on the transmission system layer for physical connectivity.

While node based network nesting is based on the concept of networks within nodes, connectivity based network layering introduces the concept of networks within networks.

For example, voice calls are carried on trunk groups and trunk groups use transmission systems. The transmission systems network layer, therefore, contains two other network layers.

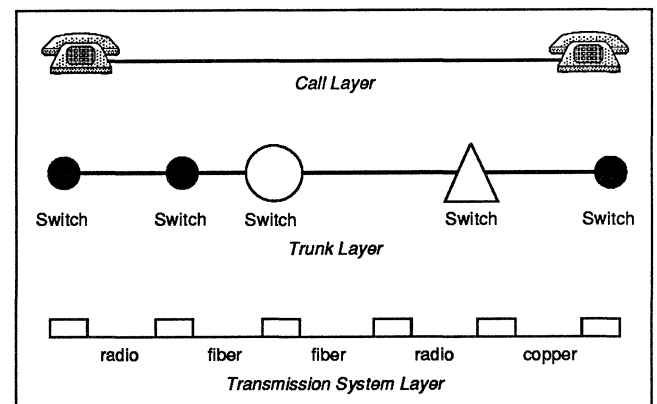


Figure 3. Connectivity based network layering.

An Interactive Network Display System for Network Management

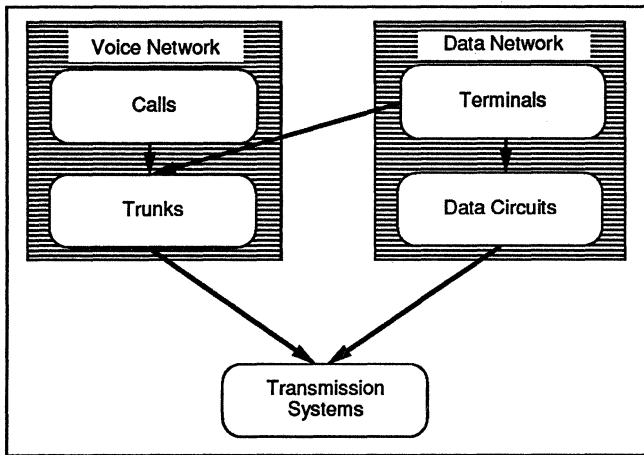


Figure 4. Network Layer—relationships.

In addition, part of a transmission system's capacity may be assigned to other independent networks. For example, a particular transmission system may be used for both voice and data networks (see Figure 4). Note that data terminals could either use the data

network or the voice network (through dial access), in which case multiple relationships exist between various network layers.

The network layer definitions are not unique and depend largely on the network being represented and the objectives of the network management system.

For example, the trunk group layer of a particular network can be further decomposed into a direct trunk group network layer, based on node type, and a tandem network layer (see Figure 5). Another example of network decomposition would be the partitioning of the network into layers based on customer network nodes and links.

The refinement of network layer definitions is a matter of choice rather than the result of fundamental network structure. Further refinement reduces the amount of information which needs to be presented at any one time, but increases the number of distinct networks which need to be independently displayed.

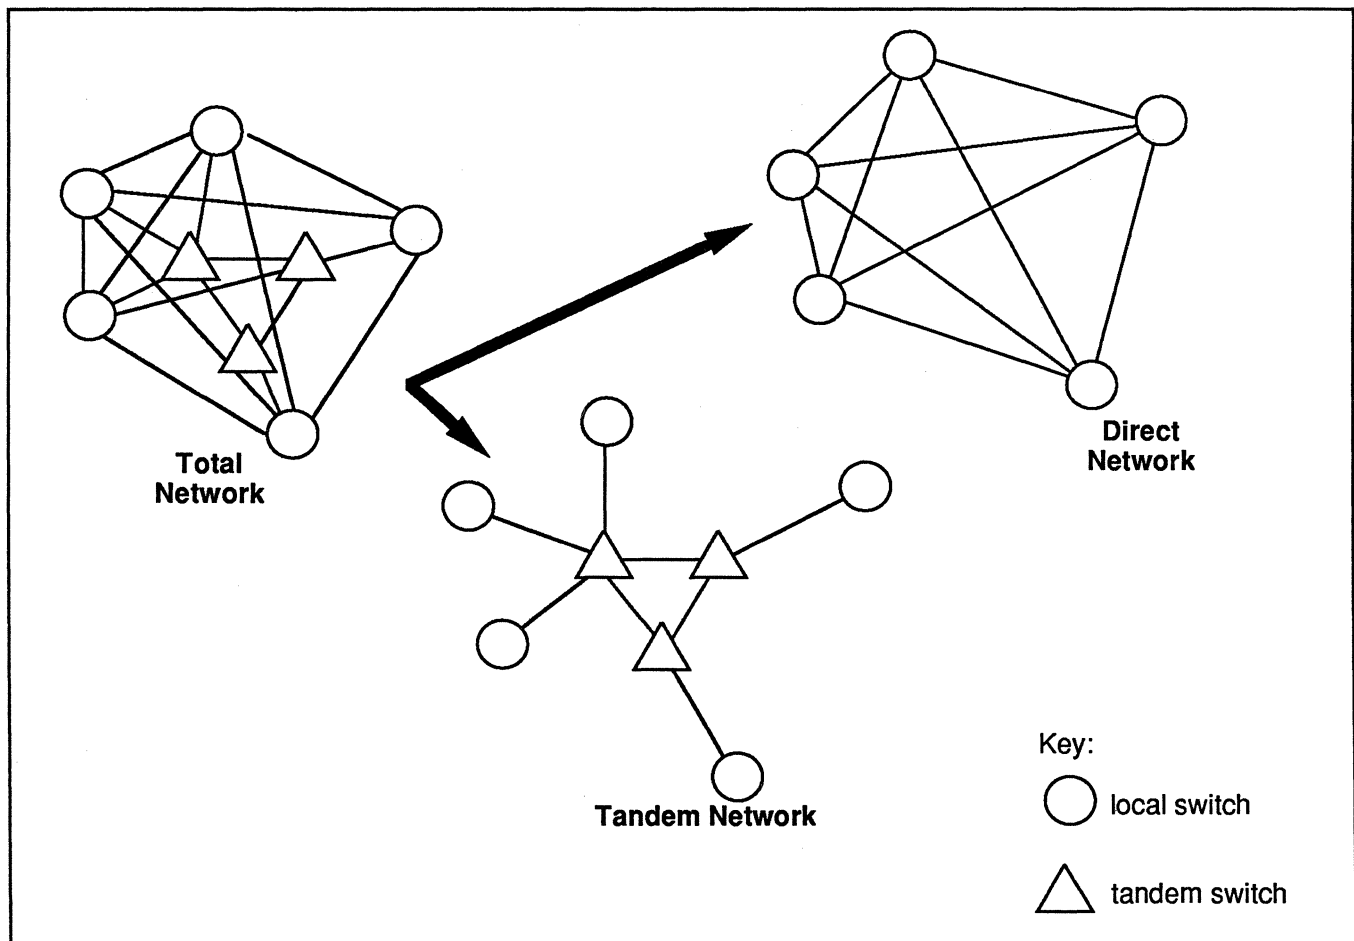


Figure 5. Network Representation—flexibility.

An Interactive Network Display System for Network Management




Figure 6. To meet network display objectives, a combination of approaches is required.

If the network is very large, further refinement of the network definition based on link type may create an excessive number of layers and may not significantly reduce the density of information in a given layer.

None of the approaches described can fully meet the objectives discussed earlier. A combination of connectivity based network layering and node based network nesting approaches, shown in Figure 6, is required (unique network representation is a subset of these).

In a combined approach, the first step is to decompose the network into appropriate layers, thus reducing the number of nodes and links being displayed at any one time.

The second step is required only if any one particular network layer contains an excessive number of nodes and links, in which case the layer is itself decomposed according to appropriate geographical or administrative boundaries.

Once such an approach is adopted, the challenge in constructing a functional network display system is in providing the system operator with the ability to access detailed information on nodes and links in any network layer in a consistent manner and the ability to explicitly examine the relationship between the various networks (each being displayed separately).

NETWORK DISPLAY SYSTEM ARCHITECTURE

An effective network management system and associated display capabilities can only be achieved through a powerful network operations architecture (see Figure 7). At the center of the architecture is the network controller, which is responsible for receiving, process-

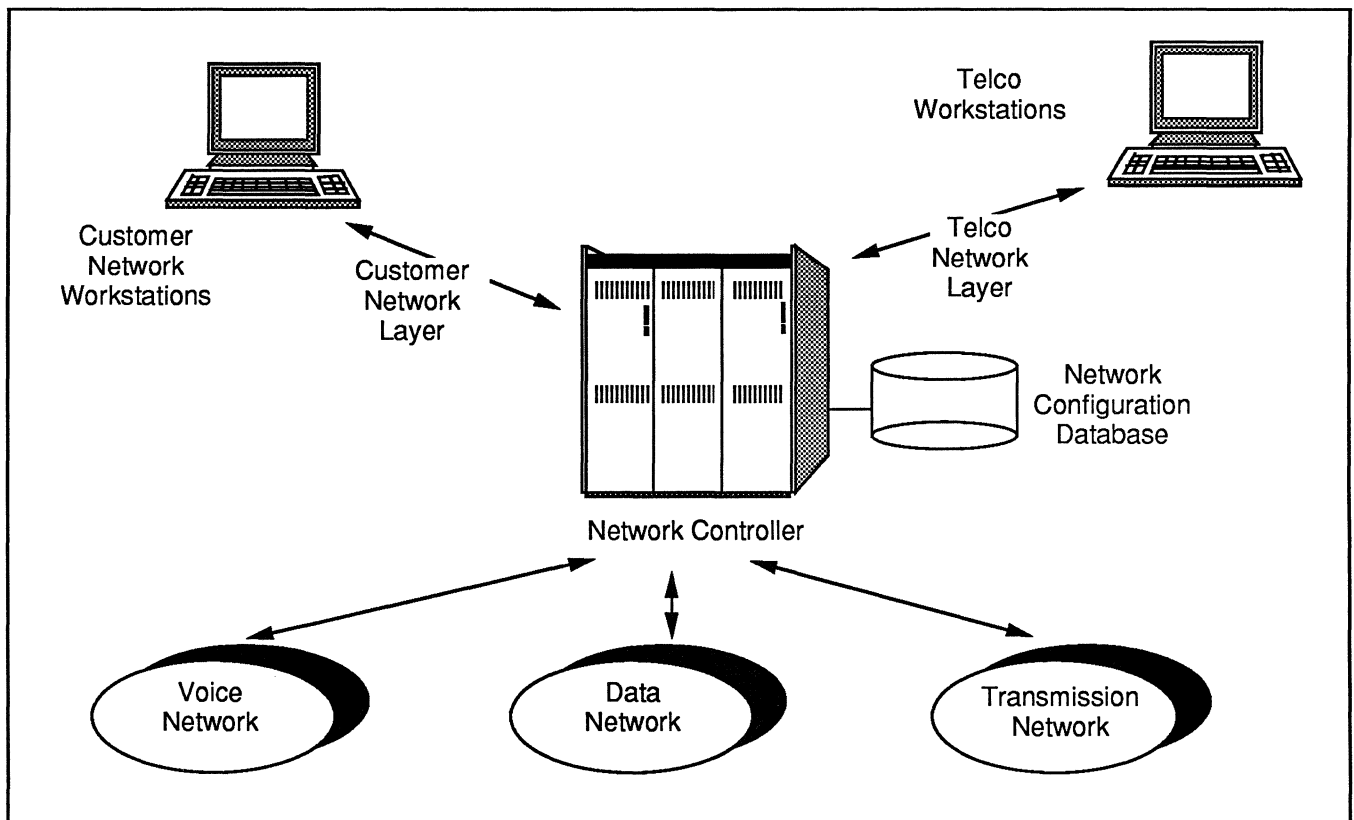


Figure 7. Network architecture.

An Interactive Network Display System for Network Management

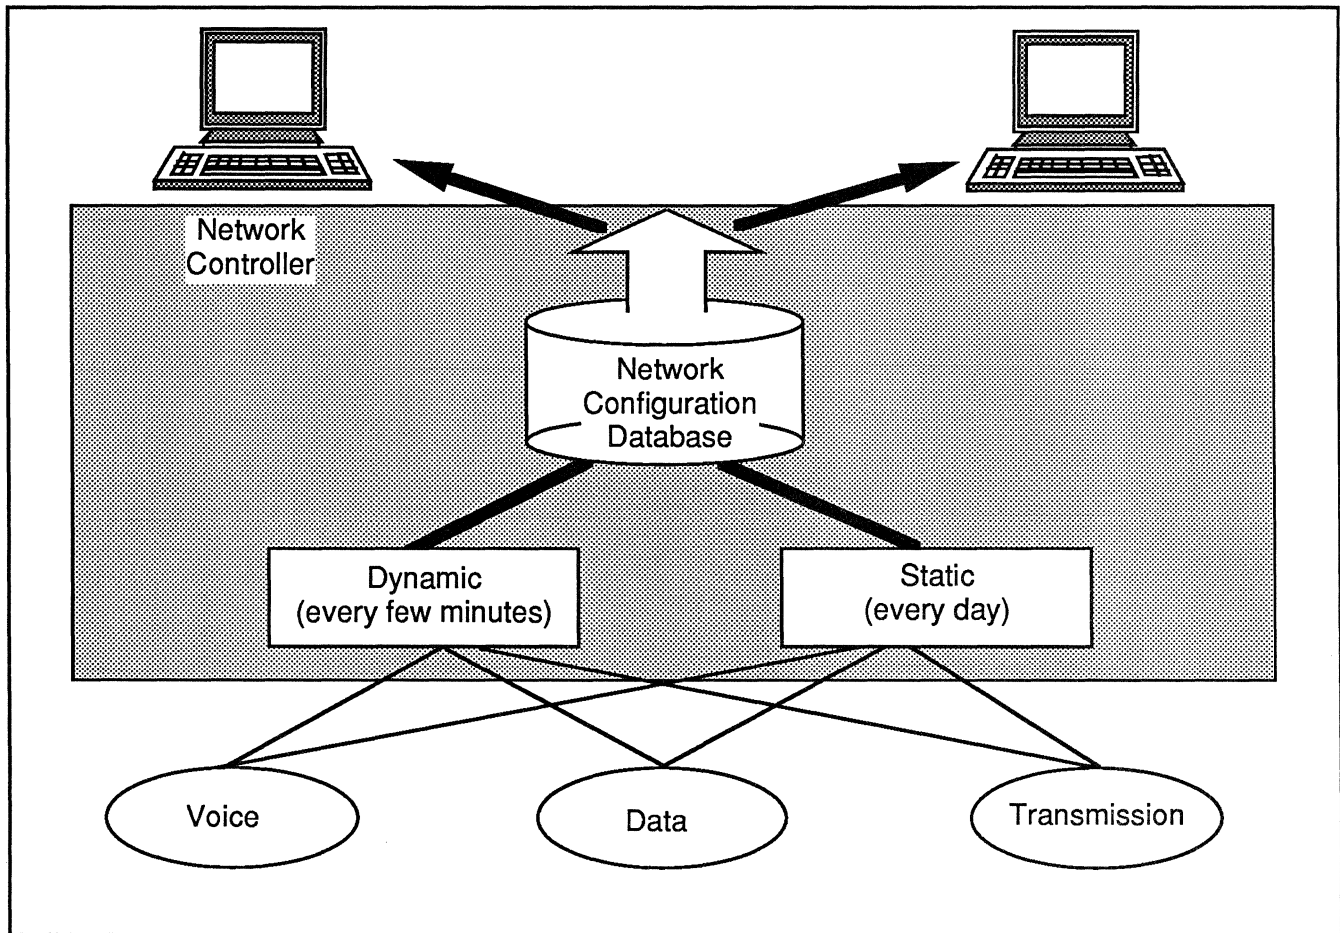


Figure 8. Incorporating dynamic and static configuration data.

ing, and storing network operational measurements and alarm information. The management of network information is through a network configuration database, responsible for the identification of the various networks and network elements being managed.

The network configuration database maintains workstation security information—workstations may receive different information in the form of different network layers. For example, customer network workstations would only receive the network layers pertaining to the network services and facilities to which they have subscribed. Telco workstations, which have access to all information, may select only those network layers which are required at any particular time; for example, a workstation may be dedicated to the surveillance of the transmission network only.

Upon reception of the network information, the workstation updates its own network database stored in RAM. This enables fast processing of near real time information on a per user basis. As part of this

update the network display is refreshed to reflect, by the use of colors, the latest status information.

To manage the vast amount of information being generated by the networks, a distinction is made between “dynamic” and “static” data (see Figure 8). Static data refers to network configuration information such as the existence of a switch in a particular location or the number of trunks in a trunk group. This configuration information changes as the result of activities such as service orders which are usually activated on a daily basis. Configuration information is always associated with one or more network layers.




Figure 9. Network Display System—implementation.

An Interactive Network Display System for Network Management

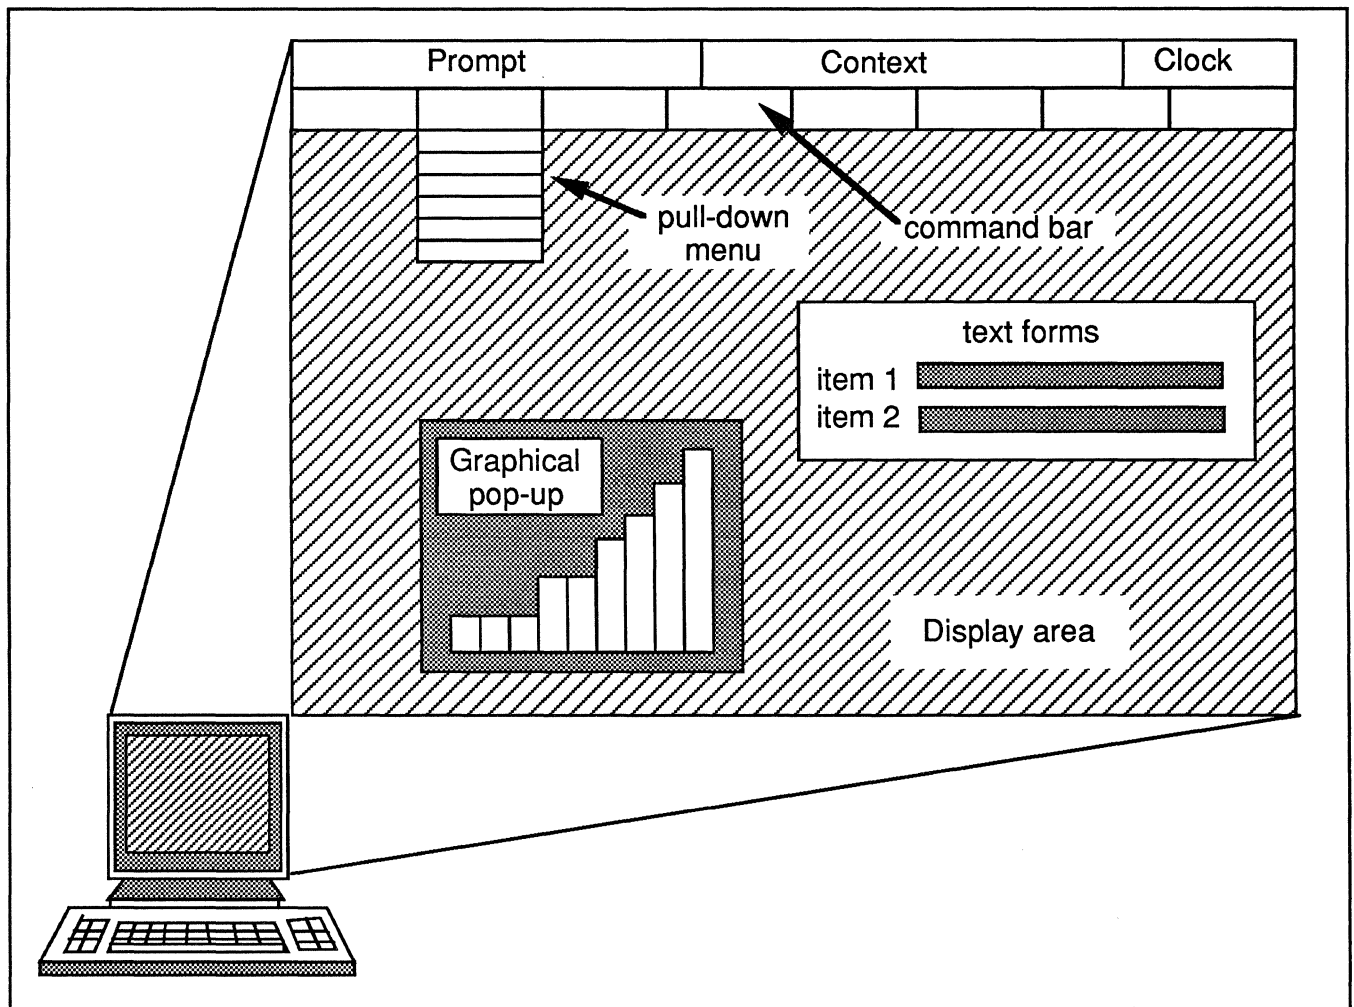


Figure 10. Network Display System—user interface.

Dynamic data refers to unscheduled events which require immediate attention, such as alarms, or scheduled events which occur frequently, such as operational measurements (every 5 to 15 minutes). Dynamic data is always associated with a network configuration element (static data).

When data is received by the network controller, either as a result of a change in the network or because of a control action by the network controller itself, the network configuration database, which contains both dynamic and static data, is updated. At the same time, data is sent to a data distributor which passes the new information to the appropriate workstations. Historical information can be retrieved from the network configuration database by the workstations on a per layer basis (provided the workstation is allowed to retrieve data associated with that layer).

The network display system (NDS) has been implemented on IBM PC AT workstations running the

XENIX operating system to allow user-oriented tasks and communication to the network controller to occur concurrently (see Figure 9). The workstation communicates with the network controller via X.25/X.400 Standard/Remote Operations, making the workstation an integral part of the network. The workstation is therefore treated as a processing element rather than as a terminal. The workstation has the ability to receive and graphically present status and alarm information, issue commands to the network controller to retrieve historical data associated with any permitted network configuration element, and issue network element control commands through terminal emulation.

The workstation is equipped with extra memory in which to store the network layers and associated static and dynamic data for which it has access. The amount of memory required depends on the number of network layers, the network size and the extent of relationship between layers. The workstation is

An Interactive Network Display System for Network Management

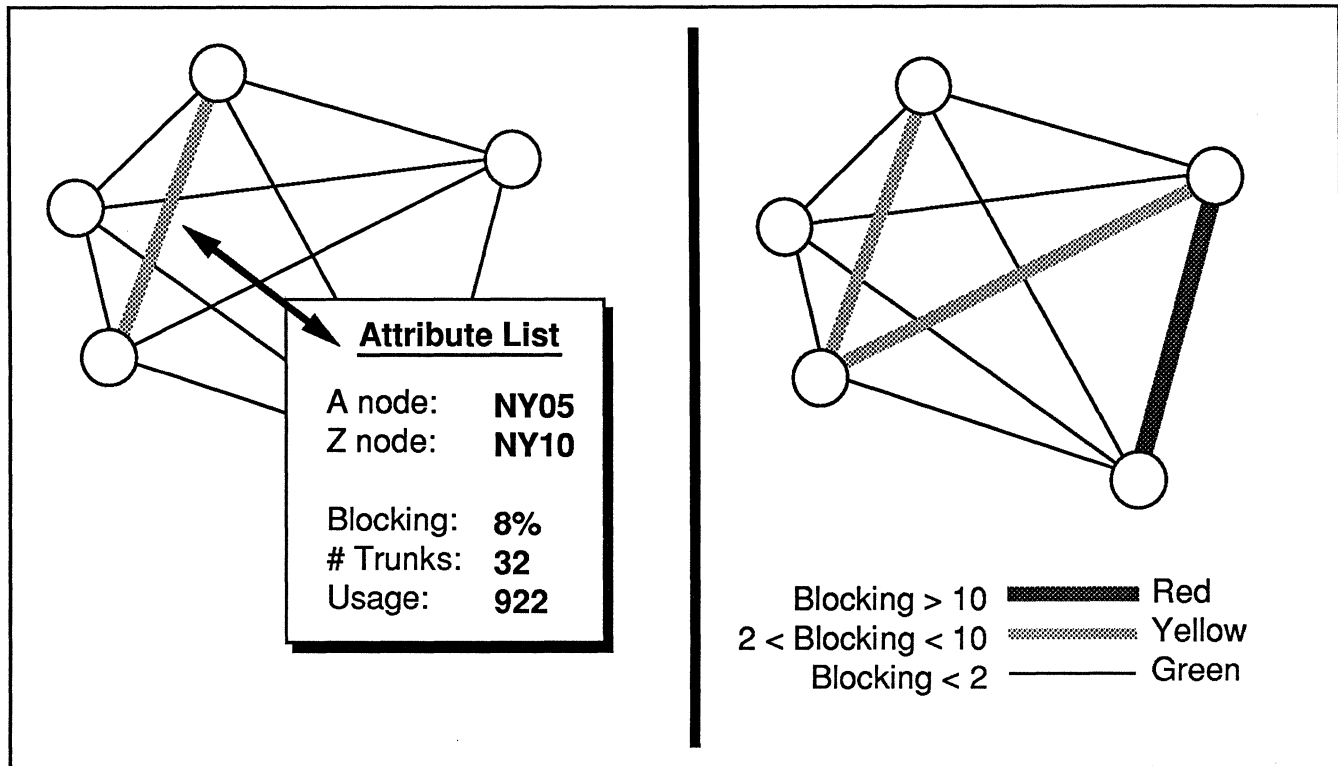


Figure 11. Sample network display.

equipped with a color monitor having a minimum resolution of 640 by 480 pixels. This resolution has been found to be adequate for single window applications. In addition the workstation is equipped with a mouse for graphically interacting with the network and a keyboard for text entry.

The workstation depends on the Business Network Management (BNM)¹ software resident in the Dynamic Network Controller (DNC) for network data collection and network management functions. BNM is responsible for interfacing to the network elements, collecting and partitioning data, and maintaining the network configuration database and associated historical information.

The Network Display System user interface has been designed for efficient interaction with the network display and uses standard screen based graphical mechanisms (see Figure 10).

The major elements of the network display window are:

1. The display area where the network and dialogs are displayed.

2. The context area which contains the title of what is being displayed in the display area, e.g., voice network.

3. The clock area which gives both the current time and the time of the last network update.

4. The command bar which is used to interact with the system.

5. Pull down menus which give the user choices within a chosen command.

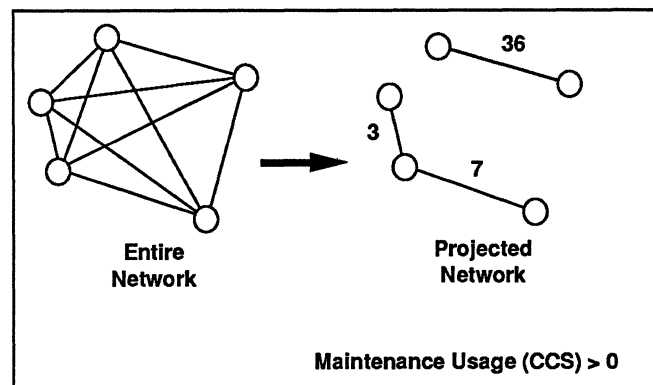


Figure 12. Network projections—example.

An Interactive Network Display System for Network Management

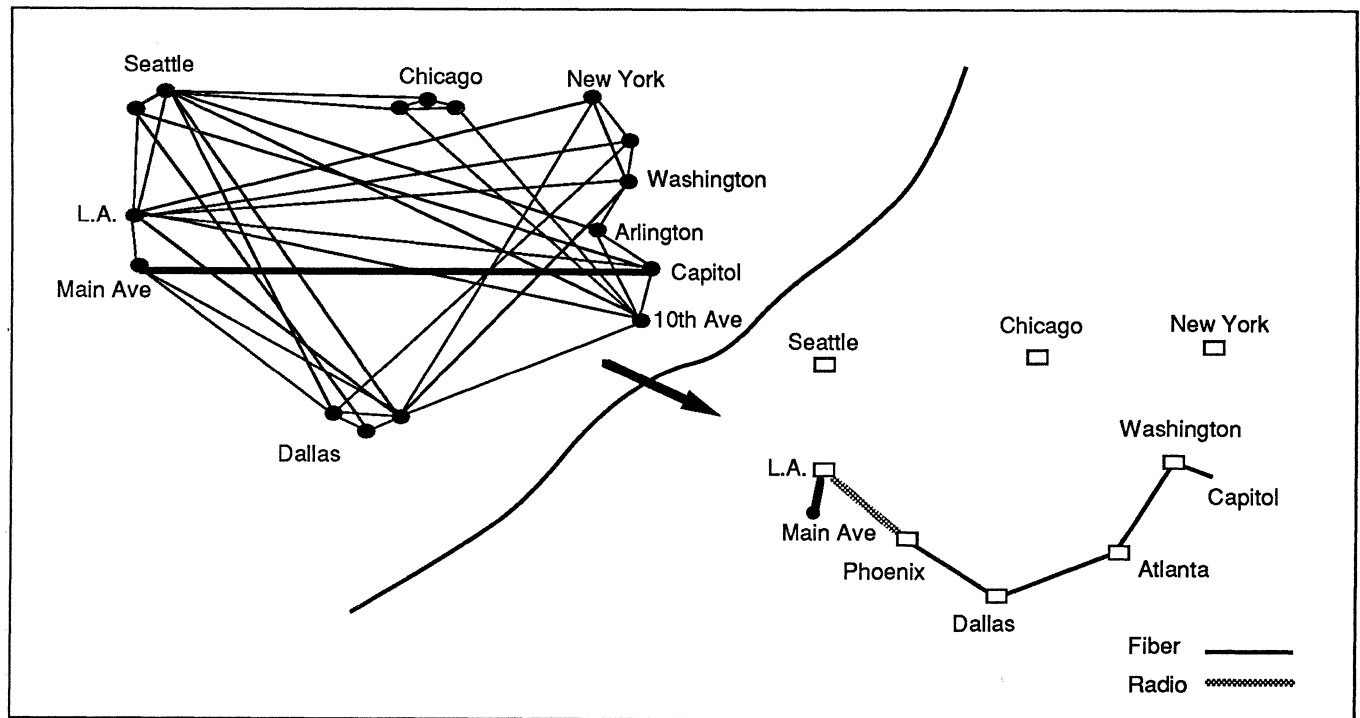


Figure 13. Network layer traversing.

6. Pop-up forms which appear when the user requests specific information or when the system needs input from the user.

7. Graphical pop-up forms are used to graphically present data.

The user interface adopts an object-oriented metaphor—upon the selection of an object on the screen (node or link) a number of command alternatives are presented to the user. The command bar usually refers to the whole network as the object.

The user initiates a session by logging in and selecting a particular network layer. The network layer appears in the display area, and the user has the ability to pick a link or node with the mouse and review the entire list of static and most recent dynamic attributes. For an historical view of a particular attribute for a particular node or link, a database command is sent to the DNC and a dataset is returned to the workstation where it can either be listed or plotted graphically.

Nodes and links are shown using colors based on user definable thresholds. The thresholds and associated colors are based on the value of dynamic attributes. For example, trunk group blocking percentage could be the attribute chosen for setting the color of links on the screen (see Figure 11). The user can define red to mean blocking in excess of 10 percent, yellow to indicate blocking greater than 2 percent but less than 10 percent and green to indicate blocking less than 2

percent. If the user selects absence of color to indicate blocking under 2 percent then only those links having blocking exceeding 2 percent would be displayed (either yellow or red depending on the actual value).

To filter the amount of information being presented, the network operator can project a subset of the network by specifying which nodes and links are to be displayed or by limiting the nodes and links to be displayed based on the value of their respective static or dynamic attributes.

For example, the user may specify the criterion that only those links for which maintenance activity has been observed be displayed (see Figure 12). The diagram shows the resulting subnetwork. As network updates are received the system will reevaluate the values of the link attributes and the display will change if maintenance activity has changed. Links may appear or disappear as maintenance usage is present or absent.

The user can also create more complex projection criteria. For example, display only those links which have more than 24 trunks and for which blocking is greater than 10 percent.

Note that similar results can be achieved by specifying colors or absence of color. However, assigning colors to more than one attribute and to specify different colors for different network layers may be confusing.

An Interactive Network Display System for Network Management

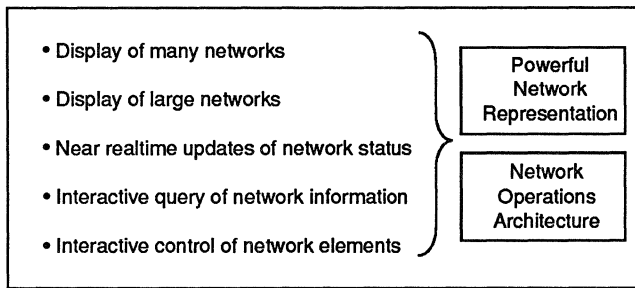


Figure 14. The objectives of an interactive network display system.

Since only one network layer can be seen at one time, a technique to graphically represent the relationship between layers is required (see Figure 13). This is accomplished by network layer traversing: the user selects a link and requests the display of dependent network layers. In the diagram, the Los Angeles Main Avenue to Washington Capitol trunk group is selected and the supporting network transmission route is displayed. The route consists in a series of fiber and radio transmission systems routed through the southern part of the country.

The selection of a link may result in more than one network layer choice. For example, the selection of a particular DS-1 may result in the dependent layers: trunk group, DS-3, and transmission system. A DS-1 may carry a number of trunk groups and be routed via a series of DS-3s and in turn the DS-3s are routed on specific transmission systems.

The ability to explicitly examine network layer relationship allows the network operator to understand network behavior. For example, if a transmission system fails, the network operator can request to see those trunk groups which will be affected.

CONCLUSION

Through the application of software and hardware technologies, the objectives of an interactive network display system have been met. (See Figure 14.)

Workstation technology, in terms of processing power, storage and graphical capabilities enables the display and user management of multiple and large networks. Advanced user interface techniques allow the network operator to browse efficiently through the various layers of the network.

Of key importance is the representation of network structures in software and in memory. This representation allows the efficient display of network status and alarm conditions. A network operations architecture, such as the one provided by the Dynamic Network Controller (DNC), allows the collection, processing, and storage of network wide information.

Most of the network display functions described in this report have been implemented and are running in the lab. The network display system is being integrated to network management applications such as BNM¹ and Dynamically Controlled Routing.² Field trials were to begin in the summer of 1988.

REFERENCES

- ¹I. Kerr and R. Coffin, "Business Network Management—New Tools for Customer Control," *IEEE NOMS88*, 1988.
- ²D. Wanamaker and A. Dorrance, "Dynamically Controlled Routing—Field Trial Experience," *IEEE NOMS88*, 1988. □

Inventory and Configuration Management

This report will help you to:

- Grasp the fundamental concepts of inventory and configuration management.
 - Add value to your existing network inventory database by using its information for performance modeling, internal order processing, and financial applications.
 - Select a network management system that provides effective inventory and configuration management features.
-
-

INVENTORY AND CONFIGURATION MANAGEMENT

Definitions

Inventory management involves keeping track of everything that *comprises* the network—hardware, software, personnel, etc. Inventory management may also include keeping track of the *absence* of something, as opposed to the presence of something, such as spare parts on controllers.

Configuration management involves keeping track of how the network is *connected*. Typically, configuration management encompasses physical connections (circuits, devices, etc.) and logical connections (sessions and applications). Configuration management may also include keeping track of how the network is connected organizationally, contracturally, spatially, and temporally.

These definitions are, by design, quite general, but generalities are necessary to support the real-world complexity of large corporate networks.

Many Networks in One

A large corporate network is actually comprised of many separate networks, including the SNA network, the electronic network, the logical network, and the voice network.

The *SNA network* starts at the systems services control point (SSCP) and travels through a series of links or devices at various levels (such as the front-end processors, the line, the controller, etc.) and ends down at a physical unit (PU) such as a terminal. All these links and devices have SNA names and are dependent on the parents in the SNA hierarchy.

The *electronic network* is the one through which electrons actually flow. This network must be continuous, or electrons cease to flow and the network has a problem.

The *logical* or the software network represents the flow of control. Typically, an application runs in a region, usually under an access method such as CICS. CICS runs under MVS, VM, or a similar operating system. Thus, the flow of control follows an established hierarchy. If any of the “parents” go down, the “children” go down within that hierarchy.

The *voice network* translates sound waves into electronic signals, then digital signals which are transmit-

Inventory and Configuration Management

ted through various components such as satellites, ground stations, microwave links, PBXs, etc. When the transmission reaches its destination, it is translated back into sound waves. All the devices that transmit the sound waves or represent them in some manner are part of the voice network.

NMS REQUIREMENTS FOR SUPPORTING INVENTORY AND CONFIGURATION MANAGEMENT

Each of the networks just described is subject to technological improvements. There is no way for a manager to predict what future technology will be incorporated into networks. Switches, minicomputers, PCs, and other now common technologies seemed like science fiction 5 or 10 years ago. Corporate networks must support all these devices and the new relationships within the network that these devices create. In addition, nondevice types such as employee directories, calendars, organization charts, etc., also require support. While these nondevice types may not be relevant to the operation today, they could be relevant tomorrow.

24 by 7 Operation

Defining and changing network components are essential activities of inventory and configuration management. These activities must occur online without taking the network management system (NMS) down. Generally, the NMS must run 24 hours a day, 7 days a week and can only be brought down for backup. Modifications must occur while the system is running.

Linking

A network's configuration needs to support linking in an M to N association. For example, a switch has M terminals connected to N computers. Suppose that any one of the M terminals can talk to any one of the N computers at any time. The NMS must be capable of representing that relationship, as opposed to only supporting a one to one association or a one to N or an N to one association. The NMS must support the fully general case of the M to N association.

Another example is representing the number of ways in which one terminal is related to other network components. One terminal can be associated with a particular controller, location, vendor, owner, servicer, software, diagnostics, problem, change, financial data, and so on. Ideally, the NMS should

represent these associations or links with the fewest possible number of keystrokes.

For example, if an NMS operator is at a terminal screen and wants to go to its controller, he/she should have the capability to move the cursor to the controller field and hit one key to obtain the management data for that controller. If the operator wants to find out what software runs on that particular terminal, he/she should have the capability to move the cursor to the software field and hit one key to display the desired information.

Aggregate Data Types

An NMS must support aggregate data types. Aggregate data is data grouped together logically. There are two basic kinds of aggregate data—array fields and structure fields. Array fields support lists of the same type of data, for example, textual comments. Structure fields support lists of different types of data, for example, a list of entries in an order.

Array field lists comprise multiple lines of the same kinds of information. Since it is not known how many lines there will be for any given array, the NMS should not require that the array length be specified up front. Trying to fill out a list of six items with only five positions on the screen can be aggravating. The NMS system needs to support arbitrary length arrays so that the arrays can be expanded to provide any needed length. This requirement should apply to structure fields as well. Structure fields are similar to arrays, but rather than listing identical kinds of information, structure fields group together different items. An address field is a practical example. The fields of an address structure include a last name, a first name, a middle initial, a street address, city, state, zip, and country. If an operator needs to create a purchase order, the NMS should provide the capability to pull up the entire structure as one entity and put it on the purchase order, rather than requiring the operator to move every field individually to the purchase order.

Mass Operations

Quite often, updates and/or changes must be made to an entire group of network components. Examples include:

- All network components in an entire building are moved to a new location.
- All prices in the inventory database require updates.

Inventory and Configuration Management

- The service contract changes on all inventoried devices.

These events require mass changes or mass updates. A mass add is one type of mass change. For instance, a mass add is required if new terminals are added to the system, all at the same location. Similarly, a mass delete is required if all terminals in one location must be pulled out of the system. And it may be necessary to do a mass print or mass unload, which allows all items to be archived together. In summary, mass operations allow the user to perform the same operation to all items on a list of records.

Reports

An NMS must support both single-file reporting and multiple-file reporting. Single-file reporting provides inventory listings. Complicated multiple-file reporting allows information to be pulled out of inventory in network order to show CPUs, front ends, lines, controllers, and drops. The multiple-file report must display all components in logical order. This requires searching multiple files and pulling out information. For example, the system first searches the CPU file, then the controller file, the terminal file, the line file, and so on in order, to pull together all the necessary information.

Reporting capabilities should support graphics, including circuit diagrams, pie charts, and bar charts. Graphics communicate information much better than text or numbers and are particularly useful in describing network conditions to upper management.

NMS reporting features should include the ability to search on any field, sort in any order, and subtotal in any way. No manager can know in advance all the reports that will be required from the NMS system. Thus, the NMS should include a generic report writer with which new reports can be built as time goes on.

Availability Tracking

An NMS must track the availability of the network components in its inventory. There are three measurements of availability—explicit, implicit, and perceived. All three have different uses and, therefore, must be tracked.

Explicit outage is the amount of time that the device itself is down. An NMS tracks explicit outage by recording mean-time-between-failure (MTBF). For example, MTBF logging may indicate that a particular device is down 1 percent of the time or 1 hour per 100 hours of running time. Once the MTBF is known,

the network manager can determine whether it is high or low relative to other similar devices in the system.

Implicit outage means that the device itself may be down or one of the device's parents is down. Implicit outage is found by tracing through the network hierarchy and viewing all the outages of all the parents, putting them together, and producing the outage of the child. If the terminal is down, it could be down because it blew a fuse—or because the terminal's line is down, or the front-end processor is down, or the host is down.

Perceived outage is the same as implicit outage except that time is logged only during the operating hours of the device, even if it went down before operating hours. If the device is only supposed to be available between 8 a.m. and 5 p.m., it does not matter if it broke down at 6 a.m. for perceived outage.

An NMS must track all three types of outage in real time. The inventory system should provide a history of outage immediately upon request on any device in the network, including perceived, implicit, and explicit, outage amounts.

ADDING VALUE TO INVENTORY AND CONFIGURATION MANAGEMENT SYSTEMS

An inventory/configuration database stores a wealth of information that can be used in numerous ways. In addition to basic tracking functions, the NMS information can be manipulated to assist in functions such as internal order processing, financial and accounting tasks, performance modeling, and call accounting.

The following sections describe ways to add value to the existing inventory database that has already been built for other reasons.

Tracking Financial and Contractual Data

As long as the NMS tracks inventory and configuration of network components, it makes sense to record financial data associated with those components. This data may include purchase information such as the date of purchase, its cost, and the tax paid. The inventory and configuration database can also record relevant depreciation information on each network component—simple financial information such as the current depreciation methods and the current book value, etc., and, for cost accounting purposes, what account or department should be charged for the item.

Inventory and Configuration Management

The NMS database can also be used to track simple contracts—the purchase, lease, or maintenance agreements—associated with each network device. Efficient tracking can alert the network manager to contracts that are nearing expiration.

The NMS database can assist with the necessary functions of chargeback, contract renewal reporting, and particularly invoice reconciliation. If the NMS tracks all contracts associated with all network devices, it can also display which invoices should be received. Those invoices can then be reconciled against the invoices actually received to see if any discrepancies exist. The NMS can even assist in simple budgeting. If there are contracts involving payment over time, the NMS can track them to see how much billing will be received in any given time period. With this data, the network manager can create a simple budget.

Internal Order Processing

One common problem in large corporate networks is controlling internal orders for equipment, such as PCs, workstations, terminals, etc. If a user wants to order a terminal, for example, the NMS should provide data on whether there are any terminals in stock (inventory). If there are, the operator can then request that terminal and receive a quote from the NMS as to the estimated delivery, lead time, and cost.

The NMS can then promote the quote into an order or hold it for a later decision. The NMS should automatically update the inventory once the terminal arrives, build a change request to install it, and also check to see if any engineering rules prohibit its installation. For example, there may not be any more spare ports on the controller, or there might be a cabling limitation, etc.

Performance Modeling

As long as there is an online database inventory system describing the network, it makes sense to run queuing theory on it to calculate the expected response times. A model of the network can be built using inventory information. After running queuing theory routines, calculated response times can be checked against actual response times. Network man-

agers can also perform a “what if” analysis—such as, if the line speed is doubled, what sort of response time will we get? If I add more terminals on this line, what am I going to see? Or, if I change the message length of this application, how is that going to change my response time? By using the existing inventory database, performance modeling can be accomplished without retyping network configuration information into another application.

Voice Call Accounting

Most phone systems provide Station Message Detail Record (SMDR) processing—a record of every call from every phone. An NMS can perform call accounting functions if it inventories all the phones in the network and keeps a record of who owns the phone. The costs associated with phone calling can be attributed to various departments, and the phone bill can be checked against what should be paid.

In addition, the NMS can process service orders. To the NMS, there is no real difference between a handset or terminal installation order.

INTEGRATION

No network management package on the market today can solve all of the problems in a large corporate network. Thus, an NMS inventory and configuration system must have the capability to interface with other systems. At the very least, an NMS inventory and configuration system must have the capability to import data from and export data to other systems and, in the future, communicate in real time with those systems via LU6.2, IBM's standard interapplication communication bridge under Systems Application Architecture (SAA).

For example, it might be necessary to communicate with a company's general ledger or fixed asset system. While general ledger is not really part of the network management system, it is still an important component to the company's data processing complex. As a participant in an overall strategy, the NMS must have the capability to communicate with other packages. □

Change Control

This report will help you to:

- Discover the difference between project management and change management.
 - Implement a change approval mechanism to manage network changes effectively.
 - Reduce network downtime by avoiding “fires” caused by improperly implemented changes.
-
-

Systematic change control is one of the least understood aspects of network management. Typically, network managers carefully analyze and plan large hardware installations and massive programming jobs. But once a network grows to more than 500 devices, a systems manager must do more than worry about project management. A manager must become adept at anticipating the constant evolutionary changes that a network undergoes.

The stage at which a company realizes that it needs to organize a change control system varies greatly, depending on the number of users and installed devices and how physically and geographically widespread the system is. When a network is small—usually under 500 devices—each user knows what the others are doing, making a change control system unnecessary. As one change manager notes, if a system administrator does not know he/she is having a problem, then there is probably no problem. All that is needed in a small network is one MIS manager who has change control responsibilities built into his/her job description.

This report was prepared exclusively for Datapro by Stewart Wolpin, a New York City-based free-lance writer who specializes in computer technology and consumer electronics. Mr. Wolpin is a former editor of *Professional Computing* and a contributor to the *Business Week* Newsletter for Information Executives.

When changes start getting out of control—when users start making independent changes that unknowingly affect each other, when communication breaks down between users and the MIS group, when one human being simply cannot handle the daily load of coordinating changes—then a new network management mechanism is called for.

Change control should not be confused with problem management or project management. Problem management is management by reaction—putting out constant system fires that increase geometrically with the number of devices in the network. Project management is the long range planning for large-scale changes, such as wiring and installing terminals in a new office complex, or replacing a database program. Although change control often affects users throughout a network, changes are still usually isolated incidents; change management does not involve critical path dependencies, the key features of project management. (See Table 1.)

The need for change control is difficult to explain to network technicians who are more interested in high technology than pedantic administration. Change control is also a difficult concept to justify financially, particularly since the payback is neither readily visible nor immediate. MIS managers must realize that building a change control system is not exciting and does not result in obvious short-term financial bene-

Change Control

PROJECT MANAGEMENT	CHANGE CONTROL
Changes with critical-path dependencies Large-scale changes Predicted or planned changes Complex, long-term changes	Changes oriented towards system/network upkeep Isolated changes that impact the entire network Changes that are anticipated but not always predicted or planned in advance Repetitive changes
EXAMPLES: Installing 50 terminals in a new building Replacing a database program	EXAMPLES: Moving a terminal Adding one terminal for a new employee Revising a report structure

Table 1. The change characteristics listed in this table provide generalized guidelines for determining whether changes fall under project management or change control. Each organization must develop its own standards specific for classifying changes.

fits. There is no easily identifiable payback on personnel or equipment versus increased efficiency measured in dollars.

The best reason to consider organizing change control is to decrease the number of network fires that need to be put out. Changes that are anticipated and implemented in a timely fashion eliminate the possibility of network-wide shutdowns. Changes themselves are not predictable, and not all changes can be controlled. But the inevitability of change is predictable; therefore, changes can be anticipated and controlled. Resources will be better utilized; in the long run, this will produce a dollar savings. If there is any lesson to be gleaned from data processing history, it is that each MIS dollar has a bigger impact if spent first on management, and second on technology.

DEFINING "CHANGE"

Change is the continuing process of fine tunings including implementing seemingly minor updates and improvements in connections, wiring, terminals, and other equipment, as well as applications software, systems programming, communications, and configuration. As any network manager knows, these characteristics are synergistic, and the relationships become more complex as the network grows. This synergism makes each individual change more critical to the overall efficiency of a network.

What is defined as a "change"—as opposed to a problem or project—will vary from company to company. Every day, minute "changes" are made within a system. A secretary adds a new record to a database file. A salesperson inputs a new order. A production manager updates inventory files. These are day-to-day data changes that affect only specific programs and users and are a normal part of everyday operations.

On the other hand, there are large-scale projects, such as adding another processor on-line or translating data into a new program. Such changes fall into the

category of project management, complete with critical path dependencies and separate software all its own.

Change control is oriented more towards system upkeep. Examples include installing a new terminal, installing a new module or modifying an existing software application, and revising a report structure. Each of these are internal organization changes that impact the overall system. In operations with thousands of terminals and tens of thousands of employees, there can be hundreds of these types of changes a week—hence the need for strict controls.

These changes may be grouped into the following three categories:

- **Application Software Changes**—changes made to the software needed to support various business functions, such as an on-line banking deposit system in a bank or an order processing system for a sales organization. The changes could be revisions in the software itself—adding a new data field or installing a new module—as well as changes in the reports generated by the software.
- **Data Center Changes**—changes often made in the computer room that can be either hardware or software related. These are changes that affect overall system operation, such as installing a direct access storage device (DASD).
- **Network Changes**—these affect the communication lines between the computer room and the individual end user. These changes can include hooking up new users, new telephone lines, new network hardware or software, such as a modem, anywhere between computer room and the remote location.

These changes are simply examples. Each network manager must identify changes that are endemic to an organization's network operation.

Change Control

MANAGING CHANGE

Typical change management begins with the understanding that change control is a crucial piece of overall system management. In many networks, a hierarchical management structure needs to be imposed on change control, complete with appropriate administrative personnel.

Change is a word that can have a narrow or wide definition—how extensive does a change have to be to fall under the authority of this hierarchical structure? Is there a definable scope for a change before it is made part of a change control system?

The likelihood of a change causing network-wide problems determines the size of the change that needs approval. One change manager likens the determinant for change to building fire codes. The fire codes are the corporate standards designed to catch the most likely causes of fires. Any changes that are made to the structure—a new room addition, central heating system, a wall knocked down—must adhere to the fire codes and be approved by the fire inspector. Conversely, there are smaller changes in a building, such as a new chair or painting a room—that really do not need formal approval.

Building changes done in accordance with fire codes are analogous to change control. Large changes that could affect the network are the changes that need managing. The small, user-specific changes that do not have network-wide impact, need not be managed as closely. Each company will have its own definitions, or codes, defining the size of its own types of changes.

To take the analogy a step further, even with fire codes—change control—there will still be fires, i.e., problem management. But the likelihood of fires occurring lessens with strict adherence to these fire codes.

The Right Technological Solution

The best way to manage change control is to apply the right technological management solution. Most SNA network managers are familiar with and use CICS and IMS transactional software. But these systems are not designed to handle continually changing requirements. In addition, IMS and CICS are designed to handle transaction volumes much higher than required for implementing change control. Transactional programs could actually hinder the smooth operation of a change control system. A relational database is a more appropriate mechanism to use for implementing change control.

A relational database program allows the change manager to connect dissimilar files and to cut through data without extensive and time-consuming query programming. By definition, a relational database is “intuitive,” since the software can answer queries by pulling specific pieces of data from more than one set of records.

A change control system is best developed over time and not immediately imposed. Just as a network grows gradually, so does a change control system. Allowing a change control system to grow slowly makes changes easier for users and managers to understand and accept.

CHANGE APPROVAL MEETINGS

A change starts with a request—this is known as opening a change. If a change management system is run properly, the change requests come from users—anyone from secretaries to executives. This openness allows the system administrators to be in constant touch with their users. But there is an attitude amongst decision makers that opening up change requests to individual users would clutter the change control bureaucracy with trivial requests. So, most companies restrict the ability to make requests.

A change is usually submitted to either an immediate supervisor or directly to an individual or group with the authority to approve a change. This request is submitted on a form, either electronic or paper, which is standardized according to the type of change requested. For instance, a software change form is different from a hardware change request form.

The change request form provides a distinct audit trail. The forms allow an administrator to not only track individual changes through their varying completion stages, but to track overall system performance and to help justify future system-wide projects.

The change request form can be simple and straightforward. For instance, a hardware change request form may include:

- The date of the request
- The name of the user requesting a change
- The Location of the equipment affected
- The type of equipment (terminal, hardware, printer, etc.)
- The date the change needs to be completed

Change Control

REQUEST DATE:	6/20/89
REQUESTED COMPLETION DATE:	7/15/89
REQUESTED BY:	JERRY Smith
EQUIPMENT TYPE:	TERMINAL
EQUIPMENT LOCATION:	Building
REASON FOR REQUEST:	NEW Employee
TYPE OF CHANGE:	Add
AFFECTED DEPARTMENTS:	SALES, ORDER ENTRY
 <i>Jerry Smith</i>	
<hr/>	
Requestor Signature	
 <i>Maria Blalock</i>	
<hr/>	
Requestor's Supervisor Signature	
 CHANGE APPROVED?: YES	
CHANGE APPROVED/DENIED BY:	JOHN HUNTER
REASONS FOR DENIAL:	
	(if applicable)
 CHANGE COMPLETED BY: JOSEPH DAVIS	
ACTUAL COMPLETION DATE:	7/16/89

Figure 1. An example change request form.

- The managers or MIS group which need to approve the change
- The names of others or departments that would be affected by the change
- A narrative of the reasons for the change

The change request form would also have room for the varying approvals needed before implementation. If the change request form is part of a relational database, it would be joined to approval forms and to other records to be completed by those who actually implement the change. (See Figure 1.)

In cases where only implementers request changes, the form is less a request than a notification of a change, but the form maintains its importance as part of an audit trail and system historical record.

The request form goes into a queue of requested changes and, based on the nature of the change, is assigned to a list of reviewers who are responsible for its approval. Reviewers examine only those changes

that fall under their technical purview, consulting managers whose areas are affected by the change.

These reviewers simply investigate requests—to look for what one change manager calls “gotchas,” a change that could bring down a network. The reviewers ensure that the change makes sense within their areas of expertise and it can be integrated into the overall network.

In a large system, there are usually enough change requests that members of review groups, also called technical support groups, examine requests every day. Generally, change requests are reviewed by one individual to facilitate the process. This individual's comments—agreement or disagreement with the proposed change—is added to the change request form.

Organizing Approval Groups

There is no paradigm for defining the technical scope of approval or technical support groups. The amount and scope of authority given to a reviewer or approval group depends on how the corporate entity itself is organized. When designing a change control system, design it within the existing corporate organization. Resist the temptation to design whole new systems and hierarchies to meet the perceived need to change control. A change control system is never bigger than the organization or network it is designed to serve. A change control system designed against the grain of the existing organizational or network structure simply will not work effectively.

For instance, many companies have a direct reporting hierarchy between individual operating groups and MIS. When a change request is made by a user outside MIS, it should be directed to the operating unit managers, and then to MIS, rather than to MIS directly. A change request made within the MIS group is also passed through management channels before reaching the appropriate approval group. Going through channels allows information to be shared and enables managers to budget correctly.

The approval mechanism should be designed along technical disciplinary lines. Listed below are five technical support groups that could form the basis of a change management approval mechanism. These five groups are not universal, but almost every company will have individuals or groups with similar responsibilities.

- **System Programming** is responsible for the system software and maintains overall operating integrity from the programming side.

Change Control

- **Network Control** runs the network from the front-end processor on out, making sure all the relevant hardware and software systems are installed correctly.
- **Production Control** is the network traffic manager; this individual or group schedules and oversees the job stream.
- **Database Administration** manages all data, making sure data is available, sufficiently organized, and that the database has sufficient space.
- **Auditors** maintain network security and also serve as the network's standards manager.

Each of these groups would review those changes relevant to their purview. If there is a request to install a disk pack, then Systems Programming and Database Administration would approve the request. If there is a request to build a new report or adjust an older report, then Auditors, Production Control, and Database Administration would approve the request. If there is a request for a new remote terminal, Network Control and Systems Programming approve the request. (See Table 2.)

Change Control Meetings

The most common change approval methodology is the change control meeting. Depending on the network's size, anywhere from 2 to 30 change approvers gather on a regular basis and review the entire list of change requests submitted since the previous meeting.

In some smaller networks, a regular change control meeting is the most efficient method of maintaining communication between MIS groups and of effectively coordinating the large number of requests and changes that require multiple approvals and authorizations.

The effectiveness of these meetings is dependent on organization. One company uses a calendarized agenda; each change request is listed in a centralized database file. MIS managers consult the list to see how many, if any, changes need decisions. The final calendar is either mailed or distributed electronically to the varying MIS managers, and the meeting is then attended by managers who need to input on specific changes.

In a large network, though, the main problem with any change control meeting is maintaining attentiveness. Most change requests have little or no relevance to the majority of attendees. The hardware specialists will not understand much of the software programming changes requested, and vice versa.

In larger organizations, regular change control meetings are a waste of resources. Change control meetings require highly trained and highly paid technical specialists to wait for the one or two items on the agenda requiring their knowledge or input.

The shortcomings of weekly meetings often are not recognized until the meetings become impossible to organize. The managers involved need time to implement the changes from the previous meeting. In practice, organizations need to start looking for an alternative to the change control meeting when meeting attendance reaches 20 or more.

Change control for a computerized network should be automated and distributed like the data that is being managed. "Meetings" become electronic and distributed; a manager can take whatever time needed to review the changes and make authorizations directly on the original electronic change request form. Change requests can be partitioned into categories, and additional communications between MIS groups can be conducted over the phone.

TYPE OF CHANGE	APPROVERS/AREAS OF CONCERN
1) Install a disk pack	Systems Programming/define it on the system Database Administration/interface to database /room left on database
2) Revise an existing report	Production Control/effect on job stream Auditors/ensure that revisions adhere to standards /effect on network security
3) Add a remote terminal	Database Administration/assess effects on database Systems Programming/are there adequate parts Network Control/are terminals available

Table 2. Approval groups review and approve changes relevant to their areas of expertise. This table lists three types of changes and the approval group members responsible for approving those changes.

Change Control

IMPLEMENTING CHANGES

Once the change request is approved, the change goes to a coordinator who implements it. In many companies, "change coordinator" is a new position and, since the job is less technical than administrative, it does not have a typical MIS job description. As an organism grows more complex, it develops organs it did not need before.

The closest analogy to a change coordinator is that of a traffic manager in an advertising firm. The traffic manager does not have the skills to write copy, take photographs, draw illustrations, or design a marketing program; however, the traffic manager knows where all the pieces are at any time during the production process.

Each change is composed of tasks, or individual work units. If a new terminal needs to be installed, there is a wiring "task" handled by electricians and an installation "task" handled by the systems programmer. Each of these tasks is performed by an implementer and overseen by the change coordinator.

Keeping track of tasks is not complicated. The relational change control software that has tracked the request through the varying approval stages now tracks each task as it is completed. The change coordinator keeps track of changes via reports generated by the change control software.

Many changes become standard, and are implemented over and over. Terminals are constantly being installed or moved, software is continually updated, and report parameters are always being revised. These regular changes are not universal, however; there is more than one way of installing a terminal. Each organization will adhere to different configurations and standards.

After a period of time, change coordinators will have compiled a library of regular changes with attendant task checklists. Once a task is completed, it is checked off the list and the next task started. This ensures a smooth implementation of most regular changes. The change library will eventually incorporate all your company standards into each change and task prototype.

A change control checklist, however, is more complex than just a mere list of items to be checked when completed. A checklist simply names a task; a change control checklist defines, sets parameters, and establishes the sequence of a task. In the case of a new terminal, for instance, the checklist would indicate that the implementer needs to assign an address to

the terminal and maps out the address assigning procedure, rather than list a simple "assign an address" instruction.

Each time an individual task is completed, it becomes a prototype—the usual way that a particular task is performed, under a specific set of circumstances. A cable needs to be pulled is one regular task for a particular change, but there are differences in the way that cable is pulled, depending on the building, floor, or terminal location, for instance. Each method of cable pulling becomes a task prototype. These task prototypes are collected into a library so that each prototype that can be applied, modified, or customized for future tasks.

The main by-product of a change library is the crystallization of the corporation's standards which govern change.

Change control keeps an audit trail of all changes and tasks throughout the implementation process, from the initial change request form to final approvals.

SOURCE CONTROL

The most common changes are those made to applications software. Large companies run hundreds of different application programs, and each is constantly upgraded and revised. There is a particular methodology that programmers use to update software—the most common type of change control—although most programmers do not view it as such. Programmers have adopted three environments:

- **User Environment.** A programmer codes and bench tests the changes on his/her own dataset.
- **Test Environment.** The production environment is simulated to test the changes before full implementation.
- **Production Environment.** This is the environment in which normal business operations run.

One project could involve the movement of hundreds of discrete pieces. What looks to be the simple addition of an entry field of a screen actually involves changing the basic structure of the database itself, compiling and linking several programs, and shifting the existing data to fit this new structure. If the dependencies of these individual tasks are not recognized, or a single finger check occurs, the entire testing process is delayed.

A bigger problem for programmers is that small mistakes are inevitable in large reprogramming jobs. The

Change Control

farther along in the testing that a bug is located, the longer it takes for a programmer to recover, since the software must be reprogrammed and the testing begun anew.

The inevitability of mistakes, as well as the predictable dependencies of individual tasks and the frequency of changes, make software changes a prime candidate for source control.

It is technically feasible to perform source control for an organization's application software using a tree-driven source control capable of understanding software change dependencies. This source control can automate these dependent changes, compile the new program, perform the link editing, and shepherd it more quickly and efficiently through the testing phase. □

Performance Modeling: Analysis of Digital Communication Systems

This report will help you to:

- Evaluate the negative effects multiplexers have on response time and other network performance aspects.
 - Grasp the functional relationships between response time and key communication parameters.
 - Understand the effects of different design criteria on response time and line loading.
-
-

INTRODUCTION

Analyzing digital communication systems becomes more complex as this field matures. It is made more complex as network complexity increases by mixing of protocols, equipment, topologies, and applications to name just a few of the parameters.

Although each individual network element may be relatively thoroughly understood, combining the elements into a complete network and understanding the interrelationships and end-to-end network effects of these combinations is a formidable task for which there are no easy answers or quick references.

Many of the capabilities in the **Quintessential** set of network design and analysis tools are based on the capability of developing a complete analysis of a network or a portion of a network, where the term *network* refers to such elements as terminals, multiplexers, host computers, modems and their connection communication links.

Our approach to the development of this capability is based on a comprehensive analysis of the problems, review of technical literature of these subjects, and de-

velopment of programs and data bases which provide quantitative analytical modeling, incorporating results of queuing theory. This approach allows a wide range of protocols and polling schemes to be selected, and allows sufficient flexibility to support modeling multiple levels of multiplexing.

This report discusses the network analysis technical foundations upon which the **Quintessential** tools are based. The presentation tries to identify from the most general perspectives where the tools fit: what problems are addressed and what solutions are chosen. The reader is assumed to be conversant with this technical field, but not necessarily intimately familiar with queuing theory result. The bibliography at the end of this report provides references to more detailed material than can be presented in this short report.

A good starting point is to understand the relation between the performance of a communication system and the load placed on it. Communications loads are best understood from a statistical viewpoint. An example might make this clearer. Recently, when a sudden earthquake struck Newport Beach, Ca., the telephone system was incapacitated. Everyone had picked up a phone virtually at the same time.

“The system was not designed to function like that,” the company representative was quoted as saying. The capacity of the system was designed to handle only the average, peak loads.

This Datapro report is based on “Analysis of Digital Communication Systems,” by Gary D. Shilling and Philip C. Miller, PhD., Quintessential Solutions, Inc. © 1988, 1989, Quintessential Solutions, Inc.

Performance Modeling: Analysis of Digital Communication Systems

Traffic in a communication system is often only known in terms of average values and distributions. Besides the mean values, there are other statistics that can be very helpful. These relate to the overall distribution of traffic variables. For example, the average number of messages per unit time may be some constant value, but the distribution of these messages may be highly clumped. Most of the messages may come at a particular time and for the remainder, the system is unused.

Load (or utilization) is the term used to describe the use made of system resources. To describe the load on a line, the algebraic relation is:

$$p = n * s$$

where p = utilization
 n = number of messages/sec
 s = average time to service a message

Utilization factor is a term to describe the relative use compared to the total capacity. Calculating the system load is the first step in analysis. Information on message types, and the pattern of their generation over time and for many devices, may then be required.

Messages are often of a *bursty* nature. This is particularly true of voice transmissions, but also may characterize terminal input. Where it is reasonable to assume a very large number of messages will be transmitted, a Poisson distribution provides a good fit. By its use, an estimate of the expected number of messages to arrive for any given time can be made. Service times are often a direct function of message lengths, therefore the distribution of the size of messages is also important. For the terminal component, the analysis model uses mean and distribution information.

The term *variance* refers to the spread or breadth of a probability distribution. With high variance, the distribution varies over a broad range. Combining this measure, with a model known as the normal distributions, allows confidence intervals to be calculated. By specifying mean and variance for the input messages, the input parameters for the terminal-level model are complete. The inputs for message length are the mean and variance using a normal distribution. Inputs for message arrivals are the mean of Poisson distribution.

The response of a system to the load placed on it can be a complex function. For example, load factors may vary temporally. The dependence of load on time may be cyclic as in daily variation or may involve lagged terms that depend on the history of the system. In most models, however, an assumption of independence or orthogonality of the input processes is often valid, where many unrelated and geographically separated




Figure 1. Three intuitive concepts of system response time. Curve (a) depicts response time as a constant, not dependent upon system capacity. In curve (b), response time is very long and changes suddenly with increased load. This corresponds to conditions at the limit—system bottlenecks, breakdowns, etc. The response time depicted by curve (c) varies linearly with load. Response time is twice as slow when the number of processes per second is doubled.

operations are taking place. The load information is the input or arguments for models that are used to predict a system's performance.

Performance is often measured in terms of a value called throughput. Throughput is a metric for the number of processes that can be carried out per unit time. Its inverse is called the response time. Response time is the measure that the system designer wants to know. There are three intuitive concepts of system responses (see Figure 1).

1. Response is a constant. It does not depend on system capacity. This is consistent with the assumption of essentially infinite capacity, or of very low loads.
2. Response time is very long and changes suddenly with increased load. This corresponds to the conditions at the limit—system bottlenecks, breakdowns, what occurs when resources are very limited compared to process requirements.
3. Response time varies linearly with load. This means that with twice the number of processes per second, the responses would be two times as slow and correspondingly, throughput would be lowered by one-half.

Devices which are used in digital communications complicate this picture beyond the range of these models. Hardware such as terminals with buffers, CPUs with complex i/o components, multiplexers with capacities for holding and buffering messages are all instances in which one feature is of particular importance. This feature is the creation of waiting lines or queues for services.

Performance Modeling: Analysis of Digital Communication Systems

Queuing Theory

Queues are familiar from everyday experience in stores, restaurants and in traffic. Having been the subject of much study, they are classified by the nomenclature:

a/b/n

where **a**= the arrival process by which messages are generated,

b= the queue discipline, describing how messages are handled, and

n= the number of servers.

For the terminal nodes of communication systems, the queue type that is most appropriate is the M/G/1. This means that the process by which message arrivals are modeled is Markovian. The service time for processing message is treated in a general format, requiring the specification of only the mean and variance. There is one server. Markov processes are ones in which the history of the system is unimportant. Message rates are modeled as being a function only of the immediately previous time period. The prediction of message rates at any time *t*, depends only on the state of the system at time *t*-1. This model is appropriate for large systems consisting of independently operated terminals.

The M/G/1 solution for queuing delay is given by the Pollaczek-Khinchin (PK) formula:

$$d = m + l \cdot m^2 / [2(1 - p)]$$

where **d**= the total delay

m= average time to transmit a message

l= number of arriving messages/second

p= utilization of the facility, expressed as a ratio.

The total delay at the terminal is composed of two elements. The first is the time to transmit the message and to handle the terminal polling overhead. The second component of the formula is the diminishment of throughput caused by the message waiting in the terminal buffer. The qualitative nature of the expression is:

- at low loads i.e., *m* near zero, the queue delay is not significant;
- at loads where utilization nears 100%, i.e., *p* near 1, the second part of the expression goes to infinity;




Figure 2. This curve depicts the total delay at the terminal. At low loads, the queue delay is not significant. At intermediate loads, a curvilinear relation between response and load exists. At loads where utilization nears 100 percent, the delay approaches infinity.

- at intermediate loads a curvilinear relation between response and load exists. (See Figure 2.)

Two important features of the queuing theory are realism and flexibility. This is brought out in comparison with related types of systems modeling. For example, a simple compartment model captures some of the delays due to queues. Consider a variable *X* whose quantity changes as a function of its current state:

$$dX/dt = -a \cdot X$$

By incorporating a process that redirects some of the output to the input in the next time period an analog to queuing can be realized:

$$dX/dt = -a \cdot X + b$$

The interpretation of **b** is material retained by compartment *X* until the next time interval, analogous to a queue. The great variation, however, in communication systems requires many unique equations.

Other approaches include building complex Markov models, with difference equations. Both methods suffer from a lack of generality. Although generic, queuing models have high predictive accuracy for the performance of buffered communication systems. In contrast to the solution of models through simulation, queuing algorithms are also highly adaptable. The algorithms may be readily applied to new problems.

Simulations may require reprogramming, and may develop problems in obtaining accurate numeric solutions. Finally, queuing theory is a mature field with an extensive literature devoted to its application to the changing requirements of digital communications.

While queuing theory relates to general systems behavior, it must be augmented by the specific details of the mechanics of interchange in communications. The setting up of a link, assuring error-free transmission, and

Performance Modeling: Analysis of Digital Communication Systems

direction of messages are the function of precisely defined protocols. In the next section, commonly used protocols are described in the context of an integrated approach.

Protocol

This section discusses general protocol concepts, protocols used in the analytic queuing model and application of these models.

Protocol modeling for hierarchical networks involves modeling the lowest two layers of the OSI model: the data link layer and the physical layer. Note that restricting the model to hierarchical networks allows only primary to secondary station communication. Peer-to-peer communication is not allowed. Consequently, the topological configurations are constrained to configurations where there is only one primary station and one or more secondary stations. A point-to-point link refers to a configuration when there is only one secondary station sharing a port to the primary station while a multipoint link refers to a configuration when there is more than one secondary station.

Each station operates link communication in one of three duplexes: simplex, half-duplex and full-duplex. In simplex transmission, data flow is in one direction only. This form is not generally used and is not modeled explicitly. Do not confuse simplex duplexity with receive-only or send-only traffic which is independent of the link duplexity.

A half-duplex link can send and receive both directions, but not concurrently. Two-way alternate is another name used to describe this mode. It is convenient in modeling a communication cycle to consider the effect of primary station processing, or host processing, to a half-duplex link. When the primary station receives data from the secondary station, the half-duplex link can be held until the primary station has completed processing the data and is ready to turn the link around to send the result of the processing. This is termed pure half-duplex.

Interleaved half-duplex refers to a scheme when the half-duplex link is not held by the primary station while the primary processes the input data. In this case, the primary frees the link for another possible communication until the response to the input is available.

A full-duplex link can support sending and receiving concurrently in both directions. Two-way simultaneous is another name used to describe this mode. Several combinations of topology, point-to-point or multipoint, and duplexity are possible. (See Table 1.)

Topology	Primary Station	Secondary Station
Point-to-point	Full-duplex	Full-duplex
Point-to-point	Interleaved half-duplex	Half-duplex
Point-to-point	Half-duplex	Half-duplex
Multipoint	Full-duplex	Full-duplex
Multipoint	Full-duplex	Half-duplex
Multipoint	Interleaved half-duplex	Half-duplex
Multipoint	Half-duplex	Half-duplex

Table 1. Possible combinations of topology and duplexity.

In order to further reduce the number of special cases that need to be considered, point-to-point is considered as a special case of multipoint.

Most line disciplines can be described as some variant of a poll and select protocol:

Poll Primary invites a designated secondary to send data.

Select Primary informs a designated secondary that data from the primary are coming.

The interaction between primary and secondary stations during a polling sequence and during a select sequence is described in Figure 3..

In general the communication protocol can be modeled by application of the following parameters. Note that it is a notational convenience to introduce the concept of input and output: input refers to data sent from the secondary station to the primary station and output refers to data sent from the primary station to the secondary station.

BPC Number of bits per character.

FRMSIZE Number of characters in a frame.

ACKFRM Number of frames per acknowledgement.

POLL Number of characters in polling message.

NAK Number of characters in negative acknowledgement message.

IF Number of characters for frame overhead for a frame sent from the secondary station to the primary station.

OF Number of characters for frame overhead for a frame sent from the primary station to the secondary station.

Performance Modeling: Analysis of Digital Communication Systems

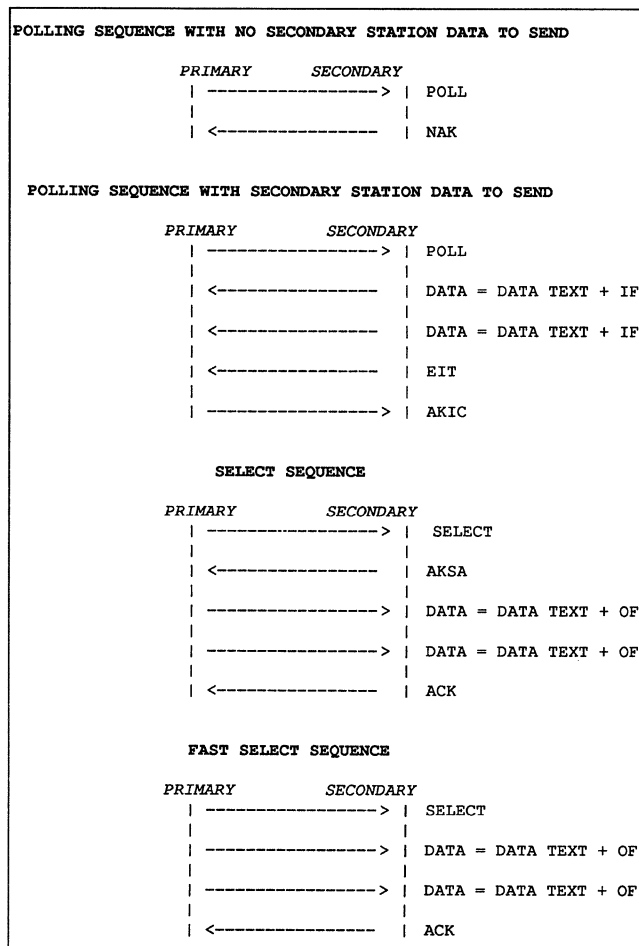


Figure 3. The interaction between primary and secondary stations during a polling sequence, and during a select sequence.

AKI Number of characters in acknowledgement message sent by the primary station to the secondary station acknowledging an input message from the secondary station.

AKO Number of characters in acknowledgement message sent by the secondary station to the primary station acknowledging an output message received from the primary station.

EIT Number of characters in the end of text message sent by the secondary station to the primary station.

EOT Number of characters in the end of text message sent by the primary station to the secondary station.

SA Number of characters in the select address message sent by the primary station.

AKSA Acknowledgement to the SA message sent by the secondary station to the primary station.

By applying the above most general protocol definition, any specific protocol can be modeled by defining as zero those elements that do not exist in the protocol being modeled. A zero value should cause no traffic impacts or delays or resource loadings in the modeling scheme.

While protocols are concerned with the details of station to station communication, the next level of organization focuses on control and allocation communication resources by the host. This is the process of polling and, in the next section, the features of polling mechanisms used in model building are described.

Polling

Improvements in line usage that yield improved price performance ratios can be achieved by various resource sharing techniques such as polling, multiplexing, line switching, packet switching, and virtual circuit allocation. Networks typically use combinations of these techniques, with some networks using all approaches to achieve high reliability and optimum price performance.

Polling is a line discipline that enables a single line to be shared among many stations, yielding a multipoint configuration. One station, the primary station, controls the line, inviting secondary stations to transmit to the primary station and sending data to the secondary station.

The standard polling configuration and technique is for all secondary stations to be connected to the same line such that all messages from the primary station are received by each secondary station and such that each secondary station sends data to the primary station over the same input line. Typically in this scheme the primary station controls use of the line by cyclically addressing each secondary station, sending output to that station and inviting that station to send input to the primary station.

An alternative scheme is referred to as hub polling, which is based on the use of full duplex lines and the capability of each secondary station listening to the transmission of data to the primary station by other secondary stations. This enables a secondary station to pass control of the input line to other secondary stations, without having to return polling control back to the primary station. During this process, the primary station can be sending data to secondary stations independent of the secondary station input process. This technique is effective in reducing polling overhead and can be quite effective when satellite facilities are used, as will be discussed later.

Performance Modeling: Analysis of Digital Communication Systems

A derivative of hub polling is the use of loop configurations, in which each secondary station acts as a repeater for the traffic that circulates in the loop. Secondary stations can insert data in the traffic and pick data from the traffic when that secondary station is the intended destination for the traffic. The success of this technique depends on the traffic delays and on the station reliability. Since each secondary station must pass on the traffic, the delay at each secondary station is cumulative to the overall message delay. Furthermore, since each station must pass on the traffic, failure of one station may cause failure of the loop and failure of two stations certainly causes failure of the loop. Experience has shown that it is difficult for loop configurations to be cost effective in wide area network designs.

Use of satellite facilities in a polling circuit introduces major delays that render the standard polling scheme unworkable because of the total delays of a full response time cycle.

Communication equipment manufacturers have circumvented this by use of a technique called poll spoofing. In poll spoofing, the primary station connects to an intelligent switch which receives the polls for all secondary stations from the primary, immediately responding to the primary according to whether the intelligent switch has or does not have data from the addressed secondary station. The intelligent switch collects data sent by the primary to addressed secondaries, responding accordingly to the primary.

At the other end of the satellite link, a complementary intelligent switch polls secondary stations as if it were a primary. If it has received secondary station data over the satellite link, that data will be sent accordingly.

The effect is that the imposition of the satellite link severs the direct connection between primary and secondary stations, spoofing each into behaving as if the direct connection was not severed.

It is obvious that techniques such as hub polling and loop configurations which significantly reduce the impact of a primary station directly polling each secondary station are effective techniques to be used with satellite connections. In each of these schemes, the primary station needs to send only a poll to one secondary station on the circuit, whereas with standard polling the primary must send an explicit poll message to each individual secondary station.

Other techniques of polling make use of such techniques as a broadcast of a message inviting any station with traffic to send that traffic, coupled with tech-

niques for avoiding collisions and recovery from collisions. These techniques are generally used only in specialized situations.

In spite of the advantages of hub polling and loop configurations, the large majority of polled circuits are operated with the standard polling scheme. Consequently, the remainder of this discussion focuses on this scheme.

The primary station is, generally, designed to poll secondary stations from a fixed polling list where each list corresponds to a specified circuit and each entry in the list corresponds to a specified secondary station. Some secondary stations can have significantly more traffic than others. Unless the polling scheme is modified to account for this imbalance, the heavy-traffic secondary stations can build up long queues and concomitant delays. Network operators have found that an effective technique that relieves the traffic imbalance is to enter the more heavy-traffic secondary stations more than once in the polling list.

An imbalance may also exist between circuits. Some circuits may be heavily loaded at other times of the day than others. In order to avoid needlessly wasting primary station resources polling low activity circuits, a pause control is sometimes inserted in the primary station polling logic. One popular implementation of the pause control is based on starting a timer when the primary station begins polling the stations in the list, and determining how much time elapsed from the start until the final station was polled. If the elapsed time is less than the value of the pause factor, the primary station delays any further activity on the circuit until the elapsed time is equal to or greater than the pause factor. The effect is that the pause control establishes a maximum rate at which secondary stations will be polled during periods of low activity. During periods of high activity, the pause control will have no effect. Consequently, the pause control feature self adjusts to the activity level.

In modeling a polled multidrop circuit which is polled with the standard polling scheme, it is important to properly account for delays caused by modems. These delays are primarily a function of the modem characteristics and the duplexity of the line. In turn, the modem characteristics are functions of the speed of transmission and the electrical characteristics of the line. Nominal modem delays range from around 10 ms. to over 50 ms. for the clear-to-send delay. Duplexity can be either full-duplex or half-duplex.

For full-duplex operation, the primary station is the only station that sends in the direction from the primary to secondary stations. Consequently, during normal operation there should be no clear-to-send modem

Performance Modeling: Analysis of Digital Communication Systems

delays in this direction. However, on the line from secondary stations to the primary station, different secondary stations alternate sending. Consequently, each time a secondary stations sends, a clear-to-send modem delay must be expected and modeled.

For half-duplex operation, clear-to-send delays for primary to secondary and secondary to primary transmissions must be modeled.

The analytic approach used for polled terminals is drawn from the class of M/G/1 queuing models. What is captured by this approach is the varying delay that is built up at the buffers under any load conditions. In addition, protocol variations are separately accounted for and participate in the calculations as additions to the average message length. The model uses two classes of parameters—mean or average values and variances (the zero and first moments, respectively). A consequence of assuming an exponential inter-arrival distribution and the corresponding Poisson message arrival distribution, is being able to treat message lengths in a general manner. This means that telling the model about the message load only requires knowledge of their mean values and their variance. The average queue delay is computed from the P-K formula, noted above. The outputs are themselves Poisson distributed and are cascaded into downstream system components. In particular, multiplexing units are frequently in place as means for linking the low-speed multidrop lines into high-speed trunk channels, offering economies of scale.

In the following section, multiplexers are described along with the models used to predict their throughput.

Multiplexer Models

Many remote devices may feed into a trunk line by means of a device called a multiplexer. Consolidating lines brings several benefits. Most important is the more economical use of high capacity trunk lines, offering economies of scale. A single T1 line, for example, has a capacity of 1.544 Mb/sec. Suppose that a terminal is being used for order entry and operates at 9.6 kb/sec. By multiplexing, more than one hundred and fifty terminals could eventually be accommodated on a single trunk line. Even greater numbers can be consolidated with an approach called statistical multiplexing. Multiplexers can also carry out performance monitoring and keep a permanent record on the effectiveness of the system.

Three basic components make up multiplexers: an input buffer, a control unit and an output buffer. Many different methods, however, are used for amalgamating digital communication lines. Analog voice chan-

nels require modulating a carrier frequency to achieve what is called frequency-division multiplexing (FDM). A group of voice channels with a basic 4 khz bandwidth can be combined into a next level line of 60 to 108 khz. Digital systems use binary pulses as the signal. Their method of multiplexing is interleaving of groups of pulses, each of which represents a separate channel. This basic approach for combining digital lines is called TDM or time division multiplexing.

The information carried from a terminal node is often highly redundant. For example, when terminal input is order entry from full menu screens, there is much repetition in the data. This repetition of terminal spacing, and control and alphanumeric coding, can be compressed by statistical means. Text information is commonly compressed by Huffman coding which takes advantage of the known statistical frequency of alphabetic character use. For graphics, run length coding is known to compress data as much as 10,000 times. Statistical methods used in combination with multiplexing can create up to 300% compression or efficiency. The hardware using this method is called STDM or statistical time division multiplexers and can handle up to three times the load of ordinary time division machines.

Queuing theory is again drawn on to provide an effective analysis of multiplexers. Available literature in the application of theory to multiplexers includes models of finite buffer size, Poisson vs non-Poisson message arrivals, and packet switching. Early work in the field has been carried out by Smith and Smith,¹ using exponential message lengths and further developed by Chu² to consider constant message lengths and multiple servers. R. Rudin³ provided an analysis of the case of limited drop lines, using binomial distributions. Many recent works have continued progress in such applications as packet switching networks.⁴

Selecting the appropriate model is an important step in applying queuing theory. The focus must be on the key factors of message processing and the effects of upstream hardware in message arrival distribution. The problem is simplified by regarding message handling procedures as operating on a character basis. Since each character takes a constant time to be transmitted and processing times are negligible in comparison with transmission times, the service time can be taken as constant.

Service time is equal to the transmission time. The relationship between the transmission time, line speed and message size is given by:

Performance Modeling: Analysis of Digital Communication Systems

$$c = b \cdot u$$

where c = transmission rate, bits/sec

b = bits/message

u = arrival rate, messages/sec

For character-based handling, message size reduces to one character. The transmission time in bits per second at 9600 bps and using 8 bit ASCII code, is .8 msec. Multiplexer processing time is faster, by an order of magnitude. The various protocols impose an overhead for addressing, error correction and control functions. The user specifies these in setup routines and they are automatically added in by the system.

The terminal is modeled as an M/G/1 system. The output of that model is Poisson distributed in terms of message frequency probability.⁴ The inputs to the mux model are therefore—constant service time and Poisson distributed arrivals. A known result from queuing theory is applied: the M/D/1 algorithm—referring to a (M) markov process, a (D) constant service time and a single server. To achieve a parsimonious model, queue buffers were not considered to be a limiting factor, that is, buffer overrides were not included as a possibility.

The average delay in an M/D/1 queue is given by:

$$m = ts \cdot (2 - p) / 2 \cdot (1 - p)$$

where m = waiting time in system

ts = average service time

p = utilization ratio

The parametric form of the solution is similar to the previous example. The numerator is associated with overhead factors- ts , the service time. The denominator is associated more directly with the traffic-dependent variable p , the ratio of load to capacity. The equations are also alike in numeric behavior. As the value of p approaches the maximum, delay goes to an asymptote at infinity. For very low loads, p is small and the delay is simply the service time, which in this case is the transmission time. Thus, the multiplexer can act as a significant processing delay to the system, but the delay is induced in a nonlinear fashion, as wait time in the queue mounts with increasing traffic. At low loads, however, multiplexer delay due to queuing is minimal. Total delay, which adds in service time, is also very small because the processing done to allocate lines takes much less time than does the character transmission time.

The model performs well given the goals of realism, generality, and computational ease. A more complex model could be envisioned which would include the possible correlation of messages. This means that the messages may not be random but arrive in chunks. Packet systems, in particular, have this feature. Feasibility of solution and the constraints of computation time also limits some of the applications. Poisson-based models were relied upon and are most tractable algorithmically. Other distributions, for example, the binomial, may involve recursive techniques for solution. Similar restrictions apply to the consideration of finite buffers where combinatorial complexity may require optimization techniques.

Thus, the multiplexer model makes a parsimonious use of queuing theory. A high degree of realism is obtained, generality superior to simulation and ease of computation. In the future, new analytic techniques and new technologies will be incorporated in the system. Some of these are considered in the next section.

Numerical Examples

Everyone has probably had a credit card checked in a store. What determines how long you will wait to get the result? This can be calculated through queueing theory. We will illustrate the principles of analysis of polled and multiplexed systems using the characteristics of this familiar transaction.

Many card readers are serviced by one host computer. When your card is checked, the clerk reads your number as encoded on the magnetic strip. In a busy convenience store, the card reader could be in almost continuous use. For this scenario, we begin by quantifying the message inputs as:

Avg. number of messages from one reader/hr. = 50
Avg. size of messages inputted from reader = 80 char.

Suppose that the host services six stores, by polling them sequentially through dedicated lines. Then the number of messages will be $6 \times 50 = 300$ messages/hour. The other information needed to define the traffic is the type of messages that are transmitted by the host. Quite often, more information is sent *from* the host than *to* it. For example, your credit card company could send not only authorization, but your credit limit and other account data. The information sent by the host is termed *output messages*. One output message will be sent for each query. The output message parameter is:

Output message size = 160 characters

Performance Modeling: Analysis of Digital Communication Systems




Figure 4. The traffic flow within a credit authorization system. One host computer serves many card readers (terminals). In a busy store, the card reader could be in almost continuous use. Quite often, more information is sent from the host than to it (such as credit limit, account data, etc.). In this diagram, all terminals must wait in one input queue. An output queue develops in the full-duplex scheme. In half duplex, the line is held until transmission from both terminals and host is completed. In full duplex and interleaved half duplex, an output queue may develop.

Schematically, the system is shown in block diagram form Figure 4.

The six terminals can all be considered to have to wait in one input queue. The average waiting time for a terminal in the queue is entered into the response time calculation. An output queue develops in the full duplex scheme. In half duplex, the line is held until transmission from both terminal and host is completed. In full duplex and interleaved half duplex, an output queue may develop.

The protocol selected is SDLC. This implies a specific organization of communication between the terminals and the host. Specifically, the host interrogates a terminal with a polling character, having the address of a specific terminal. The terminal may respond negatively (NAK) or positively to the query. An end-of-message may also be used.

Besides the protocol, full, half and interleaved half duplex lines can be specified. The variation in timing is shown in Table 2.

The cycle time is given by the following formula:

$$TC = M \cdot TP / (1 - RA \cdot TS)$$

where TC = the cycle time

M = number of terminals

TP = time for polling one terminal

RA = the poll time ratio: (TP/TS)

TS = service time

The differences between the different duplex systems are as follows: In the full duplex, the terminal is polled

and the line is held for transmission. No additional modem time is required for line turnaround. The polling in full duplex is followed by the message information without a wait for acknowledgement. The messages that are first transmitted are the output information. This is the message from the host to the terminal. The terminal transmissions are then sent or a NAK is sent if there is no message. Because the full-duplex scheme can hold output information in a queue, an additional output queue time is calculated.

Queueing theory is applied as developed for the M/G/1 queue. The model is formulated by visualizing the queue as consisting of terminals requesting to send information. All six terminals can be considered to join a waiting line as they have information to send. We are interested in how long this line will be under various load conditions. An estimate of this wait is given by the formula in Figure 5.

The components of the response time have now been calculated for a multidrop line using a standard polling scheme. In this example, we have assumed that a circuit is made up of six terminals (card readers) which are polled sequentially by the host. As the system is enlarged, line charges can become a significant expense. Often, a multiplexer is used to economically combine many slow lines into one higher speed, trunk line. In the next section, the effects of a multiplexer on response time will be calculated.

The introduction of a multiplexer changes the response times. This is because of delays introduced by the multiplexer. The delays are variable and increase with the amount of traffic. The delays are not significantly caused by the time required by processing. They are again the waiting time for retransmission that is caused by queueing in the multiplexer buffers. These buffers are necessary because the line speeds for transmission are limited and fixed. When capacities are temporarily exceeded, waiting lines for the messages are being created by placing messages in buffer storage. As previously explained, the multiplexer operates on a character basis and an M/D/1 queue model is used.

	Half Duplex	Interleaved Half Duplex	Full Duplex
Input cycle	.64	.64	0
Output cycle	1.17	1.17	1.17
NAK	0	0	.05
CPU time	.25	0	0
Total	2.06	1.86	1.2

Table 2. Variations in polling time for different types of duplexity.

Performance Modeling: Analysis of Digital Communication Systems

$$W = [N*(TP/TS) + U*(1 + TP/TS) / (1 + VS/TS**2)] / 2[1 - U(1 + TP/TS)]^2$$

where

- N= number of terminals
- TP= polling time
- TS= service time
- U= utilization ratio
- VS= variance of the service time

The results of the calculations are:

QUEUING TIME

	Half Duplex	Interleaved Half Duplex	Full Duplex
Output	0	.31	.31 sec.
Input	.590 sec.	.590 sec.	.59 sec.

Figure 5. The formula at the top calculates queuing delay when multiplexers are in use on a 4800 bps line carrying traffic from 36 terminals. The results are summarized in the table, showing the average delay at the mux for specific traffic conditions. In this situation, mux queuing delays are not a major factor in the response time. Under heavier traffic conditions, the delay will increase rapidly.

A mux is positioned in the system so that it can receive the input from six circuits. Each circuit is made up of six sets of terminals. By time allocation, the mux combines all inputs into a single, high-speed line. In other words, thirty-six of the card readers will be combined to transmit on only one higher speed line to the host computer. The method of the previous calculations remain valid, but additional wait times must be calculated for the multiplexer.

The first step in the calculation requires an estimate of the load on the lines into the mux. The transmission speed of the line to the host is selected as 4800 bps. The messages per circuit remain 300 per hour and the total input from the system will be 1800 messages per hour. The parameter of the average service time can now be calculated. To estimate the service time, the mode of operation of the mux has to be considered. A STDM works on the basis of time division, that is, the allocation of a defined portion of time to each input line. Actually the mux allocates only the amount of time needed to transmit one character. The service time is the interval needed to transmit one character multiplied by the average message rate. One further factor needs to be considered—the statistical compression that can be carried out. A compression factor of 200% is used which means that the message rate is effectively reduced by one-half. The formula for the queuing delay in this context is:

$$W = (TS/2)[(2 - P)/(1 - P)]$$

where TS= service time

P= utilization

The results give the average delay at the mux for the specific traffic conditions. The system performance is summarized in Table 3.

In this situation, the mux queuing delays are not a major factor in the response time. Under heavier traffic conditions, they can become more significant. In summary, it can be seen that the traffic from thirty-six terminals can be carried on a 4800 bps line without significant mux delays. With increases in the traffic, the delay will increase rapidly, until higher speed lines may have to be considered.

	Avg. Delay	Util.
Polling queuing	2.1 sec.	43%
Output queuing	.5 sec.	10%
Sending/Receiving	1.7 sec.	34%
CPU, I/O, modem	.4 sec.	7%
MUX	.2 sec.	5%

Table 3. The average delay at the mux for specific traffic conditions. In this example, mux queuing delays were not a major factor in the response time of 36 terminals using a 4800 bps line. With increases in traffic, however, the delay will increase rapidly.

Performance Modeling: Analysis of Digital Communication Systems

Conclusions

The analysis is being continually improved to include changes in communication systems and advances in theory. As currently implemented, each component has a corresponding queuing-based model. The outputs of the terminal model become inputs for the multiplexer. In practice, the independence of components cannot not always be assumed. Retry procedures in polling, for example, can increase the clustering of messages. The impact is on the distribution of message arrivals. Simulation results have been cited to give diverse assessments of the importance of this effect.

In some systems, sensitivity analysis highlights the importance only of mean values,⁴ while in others⁶ distributions are key factors. Theory called queuing networks is being advanced for these cascaded systems. In particular, an approach known as Jackson networks offers promise for practical implementation.

Validation has been achieved by comparative tests with alternative models. In addition, agreement with data obtained from system monitoring has been obtained. More extensive data gathering is being planned to continue testing of the model predications under all conditions of loading and configuration.

An important technological trend is towards integrated voice-data telecommunications, known as ISDN (integrated services digital network). Voice transmissions have unique characteristics—for example, they are notably *bursty*, and have silent periods of 40% during which data transmission can be interspersed. In anticipation of this technological change, a body of theoretical literature has been growing rapidly.⁷

The application of theoretical advances is a key objective in product development.

REFERENCES

Queuing Theory and Multiplexing

- ¹Smith, J.R.; Ate, J.; Smith J.L.
Loss and delay in telephone call queuing systems, Vol. 18:18-30, 1962
- ²Chu, W.W.
A study of the technique of asynchronous time-division multiplexing time sharing computer communications. Proc. 2nd Hawaii Conf. Sys. Sciences. Jan. 1969:607-610
- ³Rudin, H.
Performance of simple multiplexer-concentrators for data communication, IEEE Transactions on Communication Technology. Vol. com-19, no. 2:178-187, 1971
- ⁴Hayes, J.F.
Modeling and Analysis of Computer Communication Networks, Plenum Press, 1984
- ⁵Everling, W.
Exercises in Computer Systems Analysis, Springer-Verlag 1975
- ⁶Heffes, H.; Lucantoni, D.
A markov modulated characterization of packetized voice and data traffic and related statistical multiplexer performance, IEEE J. on Selected Areas in Communications. Vol SAC-4, No.6: 856-867, 1986
- ⁷Lee, H.
Performance analysis of statistical voice/data multiplexing with voice storage, IEEE Transactions on Communications, Vol. Comm-33, No.8:809-819, 1985

BIBLIOGRAPHY

Protocol and Polling

- Cypser, R.J.
Communications Architecture for Distributed Systems, Addison-Wesley Publishing Company
- Stallings, William
Handbook of Computer Communications Standards, Vol.1, MacMillan Publishing Company, New York

Queuing Theory and Multiplexing

- Martin, J.
Systems Analysis for Data Transmission, Prentice-Hall, 1972 □

Transport Management

This report will help you to:

- Determine how your organization can gain greater control over its communications resources by implementing transport management capabilities.
 - Evaluate routing concepts which can help your organization adapt to changes more easily.
 - Identify the product capabilities necessary for providing adequate transport management.
-
-

Users have long recognized that the integration of voice and data over private facilities offers substantial economic benefits. There are several issues that users must consider when designing and implementing such a network. These issues include:

- The quality of transmission, instead of mere bandwidth capacity.
- The integrity of the applications and information, rather than the scheme for physical delivery.
- The extension of network management and control through the public network at the channel level, not just the aggregate level over the backbone.

Without control over communications resources, companies would not be able to support internal operations, let alone position themselves to take advantage of changing tariffs, new technologies, and emerging services like ISDN. Thus, lack of control

inflates operating costs, and may render the organizational infrastructure inefficient and, in extreme cases, obsolete.

The combined advances in technology and software have produced products and capabilities which help users manage today's complex wide-area networks. These advances contribute substantially to the organization's ability to establish its strategic competitive advantage, and then maintain it over the long haul.

CIRCUIT SETUP

Transport management begins with circuit setup. The first step is to create a profile for each circuit. Users can then employ the circuit profile to select various options. These options include assigning bandwidth priority, manual routing, and indicating whether downspeeding will or will not be allowed during line failure scenarios. The user may also establish a set of qualifiers which are used to automatically route the circuits over the correct aggregate types. These qualifiers include mandatory, desirable, undesirable, not allowed, and not important. Users, for example, may choose to create an "encryption-mandatory" circuit profile for an application requiring a high degree of protection from unauthorized access.

This report was developed exclusively for Datapro by Nathan J. Muller. A former consultant, Mr. Muller has 18 years' experience in the computer and telecommunications industries. He has written extensively on all aspects of computers and communications, and is the author of "Minimum Risk Strategy for Acquiring Communications Equipment and Services" (Artech House, 1989).

Transport Management

Routing is also accomplished according to user-specified aggregate parameters. The network management system (NMS) will automatically optimize routing on specified aggregate delay, error rate, and reliability parameters. The NMS will use the shortest path between the two circuit end points that meets the requirements specified in the circuit profiles.

The ability to establish circuit profiles insures that applications are assigned to the most appropriate circuits. Beyond that, circuit profiles allow the network management system to match applications with appropriate circuits during automatic rerouting to insure that no voice and data calls get lost when bypassing failing lines.

ROUTING

The NMS determines routes by specifying the various circuits through the entire network. The operator enters route end points only. The circuits to be routed are identified by the network manager on a per circuit or user group basis. The NMS chooses the best path based on the match between the circuit profile and on the attributes and parameters of the aggregate.

The characteristics of each aggregate are determined during configuration of the network. The quality-based parameters of each aggregate include delay, error rate, availability, and user-defined attributes. The data entered is used in conjunction with circuit profiles to insure optimum routing. In this way, a route is performed based on quality considerations to insure the integrity of the applications.

By employing the delay measurement capabilities of the NMS, users can make intelligent decisions on rerouting, equipment upgrades, or choice of contention schemes to achieve terminal-to-host response time objectives. In a multidrop network, a delay measurement capability can even be used to add precision to polling, which may permit the addition of more drop locations without detracting from the overall response time of the network. Such capabilities allow companies to get the most out of their sizeable investments in networking technology.

Automatic Rerouting

An automatic reroute capability insures that the connectivity of critical applications in today's information-intensive business environment is done quickly, efficiently, and without the need for human intervention. In addition to the simple prioritized rerouting and bumping schemes commonly used, the capability to automatically downspeed circuits during

automatic rerouting insures that *all* users continue to communicate, rather than just a few.

Traditional bumping schemes resulted in service denial. Today's sophisticated NMSs automatically downspeed voice and data applications during automatic rerouting to insure uninterrupted network connectivity for all applications. For example, downspeeding voice circuits "on the fly" from 64K bps to 32K bps, or 32K bps to 16K bps, increases the number of available channels by a factor of 2:1 or 4:1, depending on the level of ADPCM compression invoked. Thus all users are permitted to stay on-line during intelligent automatic rerouting; none get bumped off. As a result, the enterprise can continue normal business operations, even though a few lines or a whole segment of the network has failed—and with little or no sacrifice in circuit quality.

This scenario offers many other advantages. Downspeeding permits continued operation with efficient T1 bandwidth fills. Efficient bandwidth fills produce cost savings. By employing the capability to downspeed during automatic rerouting, fewer T1 lines are required—offering even greater savings than priority bumping schemes.

When optimizing circuit routes, the NMS conforms to the quality-based routing parameters established during initial route determination. The system automatically calculates rerouting based on each likely failure. In the event of a failure, the system automatically recalculates optimized routing based on current network conditions. After restoration, the system will again automatically calculate the best rerouting should a second failure occur on the network. The system will then be ready to handle the next emergency until there is not enough of the network remaining through which to reroute traffic. Routes assigned in failure scenarios are always available to the network manager through retrieval from the system controller.

Time-Oriented Reconfiguration

In today's competitive business environment, high-capacity digital networks must adapt to accommodate changing application needs. Time-oriented reconfiguration allows users to do this on a scheduled basis. Circuit routing may be altered to accommodate applications that change from day to night. Since voice traffic tends to diminish after normal business hours and data traffic changes from transaction-based to batch, leading-edge NMSs provide the means to adapt automatically, without operator intervention.

Transport Management

In other cases, organizations might want to alter their networks to track the business day around the world. For example, through time-oriented reconfiguration, the circuit end points used for order entry applications are adjusted automatically to permit order entry terminals on the West Coast to come on-line as those on the East Coast shut down. This means that customers on the East Coast can still be serviced after normal business hours by terminal operators on the West Coast. Instead of losing business from dissatisfied customers, the company enhances its service image, thereby gaining a significant strategic competitive advantage.

OFF-LINE NETWORK MODELING

Network planners are now able to simulate off-line various disaster scenarios on an aggregate or node level anywhere in the network. Thus planners can test and monitor changing conditions and determine their impact on network operations.

Network control center personnel can create models that typify network failures and determine the viability of recovery strategies—all without impacting current network operations. The various models can be stored within the NMS system controller and/or printed out for hard copy reference. In this way, disaster contingency planning is based on tested, rather than “best guess,” scenarios. Failure can be simulated for aggregates, nodes, and circuits.

Through menu selections, the network planner can “create” aggregate failure. The aggregate failure will invoke the automatic rerouting mechanism, which will determine optimum channel routing from the quality-based parameters, but without downloading the changes into the network. The simulation will not affect the active configuration database.

The network planner is able to “create” a node failure by forcing total aggregate failure to that node. The automatic rerouting algorithm will reroute channels based on the entire node failing.

The network planner can also display and/or print out the circuit routing as a result of the test. Reroute scenarios can be stored and retrieved to facilitate disaster contingency planning.

PARTITIONING

The ability to segment the network for control purposes is referred to as “partitioning.” The ability to segment network management and control requires multiple network management platforms on the same

network. This makes it possible for far-flung networks to pass network control capabilities between time zones, implement management and control even if nodes become separated from the rest of the network, and continue receiving network status information from any site on the network at all times.

The concept of partitioning emerged out of the need to reconcile the twin corporate needs for both central and local control. Users recognized early on that it is impossible to share control of the various facilities and equipment *on the same network*—at least not without allocating network control to all interested parties. Obviously, if every corporate entity had an equal say in how the network is controlled, chaos would reign, bringing business operations to a standstill. Yet without local control, the entire organization’s ability to respond effectively to environmental changes would be significantly diminished.

This inability to respond is potentially disastrous, especially for acquisition-oriented companies, multinational corporations, and diversified conglomerates.

For example, the manager of an engineering division network in California requires a LAN gateway to a London subsidiary. To maximize efficiency, cost savings, and applications availability, the California user should select a route which goes through the headquarters network in Boston. This is difficult to do, however, when control is equally shared by all network locations. Thus the California user gains rights to implement any change he/she desires in the headquarters and London segments of the network—with or without the consent or knowledge of those parties. This can lead to anarchy.

To preempt this possibility, totally separate digital networks have evolved. These networks serve the needs of local users, but at the expense of the corporate need for efficient connectivity and cost control. Consequently, the need for a more logical allocation of resources for management purposes became necessary.

The technology best suited for such operating environments blends central and distributed control architectures. This hybrid architecture is “centrally weighted,” supporting one master controller which can configure the network and monitor performance. The remaining subordinate controllers have the ability to display status information and perform diagnostics. In this way, a certain amount of control is distributed among the partitioned business entities within the organization, while overall control is retained at the corporate level. Thus, local operations personnel can move, add, and change terminals; issue passwords; initiate diagnostics; and implement

Transport Management

restoral of failing links and equipment. Through the master controller, partitioning is enforced and network resources allocated in ways that benefit the entire organization. Since subordinate nodes are partitioned, there is no chance that they can disrupt the operations of other nodes, either inadvertently or willfully.

Diagnostics, status monitoring, verification functions, and off-line configuration testing are performed concurrently at all controller sites on the network. Password protection is assigned by the master operator to all locations. The ability to interrogate without changing status, and/or the ability to make temporary changes, is supported as restricted access. Further segregation of control functions by password is possible on a function-by-function basis. For example, active alarms may be sent to the primary controller only. Operators at secondary controllers may view these alarms on an inquiry basis.

Each subordinate controller is periodically updated to reflect the current state of the master database. When an aggregate trunk failure isolates a remote node, communications with the master system controller may be initiated automatically over a dial-up line via an integral modem. This process insures the integrity and continuity of network management. If the master node fails, network control is delegated to a predetermined subordinate controller, which hands back control when the master node comes back on line.

This partitioning feature insures that intelligent automatic rerouting occurs on all sections of the network that may become isolated by failures. The primary operator can delete a controller from the network for any reason by merely deleting it on the configuration screen. All other controllers are automatically notified of the change.

A GLOBAL PERSPECTIVE

The management system for each type of network component (e.g., modems, statistical multiplexers, packet switches, time division multiplexers) requires its own hardware and software, as well as specialized technicians who can interpret the alarms and reports before invoking the appropriate diagnostics and restoral actions.

Modems, for example, may have their own management system that includes the ability to monitor, test, and control a network of modems from a central location. Multiplexer and packet switch manufactur-

ers go one step further in the network hierarchy, enabling the user to manage the bandwidth delivered over high-speed lines. Computer vendors offer users the ability to monitor the performance of host-based systems through a single terminal, like IBM's NetView. Digital Equipment Corporation is encouraging third-party development of modules that will plug into its emerging Enterprise Management Architecture (EMA).

AT&T recently introduced its own network management system, ACCUMASTER Integrator. This is a central component of AT&T's OSI-based Unified Network Management Architecture (UNMA), which promises eventually to link a variety of communications devices into a single, consolidated network management system.

Each of these approaches, however, reflects the traditional strengths of the vendors. By choice or by dictum, network managers ultimately will be confronted with the task of integrating these diverse NMSs. The sheer diversity and complexity of today's sprawling networks requires something more—an integrated network management system that can view the entire network, including autonomously controlled subnetworks and hybrid networks composed of public and private elements.

Open Systems Interconnection (OSI) will eventually enable the network components of different vendors to tie into the same network management system. With alarms, performance measurements, usage statistics, and diagnostic test results in a standard format, every piece of the network can be more efficiently, monitored and controlled.

Some industry observers claim that OSI-compliant products are still three to five years away. It will take longer for users to cost-justify integrating such products into their existing networks. While users can plan for future OSI implementation, OSI does nothing to address the present requirements of users who urgently need to manage their sprawling networks more precisely, efficiently, and economically.

Rather than wait for standards, some vendors have "opened" the architecture of their NMSs by providing hooks into IBM's NetView, AT&T's UNMA, and Digital's EMA. In this way, they hope to position users to take advantage of the best mix of complementary management options: their own data communications networking expertise, plus the host-oriented management expertise of IBM and Digital, and the public network management expertise of AT&T. □

Evolving Security Management Standards

This report will help you to:

- Discover how security management is an integral part of evolving standards for network management.
 - Review essential components of the X.500 Directory and how they relate to security management.
 - Prepare for implementing OSI security measures in future network management implementations.
-
-

SECURITY MANAGEMENT

Evolving international standards for network management include security management as one of the five functional areas (along with fault management, configuration management, performance management, and accounting management.)

Security management concerns both the network and the network management system (NMS). Specifically, the NMS must be secure; it should provide tools for monitoring and controlling security (such as audit trails and key management); and it should ensure network security proper (including facilities for encipherment, digital signatures, etc.)

Security management encompasses many administrative procedures and operations needed to support and control a secure communications network. Security management functions include the creation, deletion, and control of security services and mechanisms; the distribution of security-relevant information; and the reporting of security-relevant events. It also concerns security of the network management mechanism itself.

Generally, security management includes operations that are outside of the normal communications network routine.

Entities that are subject to a single security policy and administered by a single authority are sometimes collected into a "security domain." By contrast, networked systems may fall into many domains, thus requiring different security policies. This is what makes security management in a large communications network especially complicated.

Security management standards support the control and distribution of information to various end systems that provide security services and mechanisms and that report on security services, mechanisms, and security-related events. Therefore, security management requires distribution of information to these services and mechanisms, as well as the collection of information concerning their operation. Examples are the distribution of cryptographic keys, the distribution of information on an entity's access rights, the reporting of both normal and abnormal security events (audit trails), and service activation and deactivation. However, security management does *not* address the passing of security-relevant information in products that call up specific security services, e.g., parameters in connection request.

The security-related information that open systems need conceptually constitutes a database called the Se-

This report was developed exclusively for Datapro by Daniel Minoli, an adjunct professor at New York University's Information Technology Institute. Mr. Minoli is also a full-time data communications researcher and strategic planner.

Evolving Security Management Standards

curity Management Information Base (SMIB). Each end system contains the necessary *local* information to enforce appropriate security; the SMIB holds information for the *entire* system. Security management also may require the exchange of pertinent security information between cooperating system administration processes to create or update the SMIB.

In most cases, the security-related information is exchanged over a data communications connection. Security management protocols and the communication channels carrying the management information are especially vulnerable; hence, they require protection. In other cases, the security-related information may be transferred through non-OSI communications paths, and the local open system's administrators may update the database through local, non-OSI standard methods. The SMIB concept does not presume any storage implementation or format; this database can be implemented as a centralized or distributed database.

FOUR CATEGORIES OF SECURITY MANAGEMENT

The OSI standard distinguishes four categories of security management activities: system security management, security service management, security mechanism management, and security of OSI Management.

System Security Management

System security management involves the entire open system network. It encompasses the following activities:

- Overall security policy management, including updates and maintenance consistency;
- Security-related event-handling management;
- Security audit management;
- Interaction with other OSI management functions;
- Interaction with security mechanism management; and
- Security recovery management.

Security Service Management

Security service management involves particular security services. It includes the following activities:

- Determining and assigning target protection for the service;
- Negotiating available local or remote security mechanisms;
- Invoking specific security mechanisms via the proper security mechanism management function;
- Interacting with other security service management functions and security mechanism management functions; and
- Selecting a specific security mechanism to provide the requested security service when alternatives exist.

Security Mechanism Management

Security mechanism management encompasses management of the following functions: cryptographic keys, encipherment, access control, data integrity, authentication, routing control, and notarization.

Cryptographic keys are vitally important to the organizations that use them to encrypt and decrypt information; thus, protecting their integrity is extremely important. Key management within the OSI environment involves the following:

- Generating suitable keys at intervals to satisfy security requirements.
- Determining which entities should receive a copy of each key, in accordance with access control requirements.
- Distributing the keys to the proper entities.

Encipherment is defined by the OSI as (a) interaction with cryptographic key management; (b) establishment of cryptographic parameters; and (c) cryptographic synchronization. The existence of an encipherment mechanism implies the use of key management and of common ways to reference the cryptographic algorithms. A common reference for cryptographic algorithms can be obtained through a notary (a trusted third party that assures the accuracy of its characteristics; for example, content, origin, time of creation, and delivery) or by prior agreement between the communicating entities.

Access control consists of the management and distribution of security attributes, e.g., passwords, or updates to access control or capability lists. Access control management may also encompass the use of an

Evolving Security Management Standards

access control protocol between communicating entities and other entities providing access control services.

Access control password tables are often vulnerable. The tables gather all the passwords in one data file; anyone who obtains this information can impersonate any of the system's users. To guard against this possibility, the access control manager must make sure that the password file does not contain the passwords themselves, but only *images* of the passwords under a one-way function that is easy to compute but difficult to invert. Given only the image of the password, it is very difficult to find the input string which produced it. This reduces the value of the password table to a potential intruder, since its entries are not actual passwords and, therefore, not acceptable to the password verification routine.

Data integrity management includes interaction with key management, establishment of cryptographic parameters and algorithms, and use of data integrity protocol between communicating entities.

Authentication is a security service that ensures a claimed identity is correct. It requires the distribution of descriptive information, passwords, or secret codes to entities that perform authentication. It also includes the use of a protocol between communicating entities that provide the authentication services.

Traffic padding management involves maintaining rules for communications, including such things as prespecified data rates; specifying random data rates; and specifying message characteristics such as length, and so forth.

Notarization is the distribution of information about notaries across networks, the use of a protocol between a notary and the communications entities, and interaction with notaries.

Security of OSI Management

Security of OSI management involves X.500, an emerging OSI distributed directory standard. The OSI standard assumes that a system is secure; therefore, OSI is not concerned with the host as a discrete entity, but with security of the interface between the host and the outside world. OSI security functions concern aspects of a communications path that permit open systems to achieve a secure information transfer. They do not cover measures that may be required to protect equipment, installations, and other entities. It is important to note, therefore, that additional security measures may be needed in end systems, during installations, and at the site itself.

The X.500 standard, discussed below, defines those security measures needed to protect the interface between the host and the outside world.

THE X.500 DIRECTORY IN SECURITY MANAGEMENT

OSI standards are primarily concerned with security at the interface between the host and the outside world. They specifically address aspects of a communications path that permit a secure information transfer.

The X.500 Directory is a collection of open systems that hold a logical database of information about a set of objects. A draft joint CCITT and ISO set of documents for the Directory was published in December 1988; the standards are commonly known as X.500.

Although the X.500 series of Recommendations refers to the Directory in the singular, it reflects the intention to create, through a unified name space, one logical directory composed of *many* systems and serving many applications. Whether these systems interwork depends on the needs of the applications they support. Applications dealing with nonintersecting worlds of objects may not relate. However, the single-name space facilitates later interworking should the needs change.

The X.500 Directory supports many capabilities required by applications, management processes, OSI layer entities, and telecommunications services. Among the capabilities the Directory provides are network addressing functions and the identification of legal network subscribers—important functions in security management.

The Directory also supports “user-friendly naming,” whereby objects can be referred to by names that users find easy to remember, and “name-to-address mapping,” which allows dynamic binding between objects and their locations. The latter capability allows “self-configuring,” i.e., removal and the changes of object location do not affect OSI network operation.

Information held in the Directory is collectively known as the Directory Information Base (DIB). The DIB stores user credentials electronically. Recommendation X.501 defines the DIB and its structure. The DIB is made up of the information about objects. It is composed of (directory) entries, each consisting of a collection of information on one object. An entry is made up of attributes, which have a type and one or more values. The types of attributes present in a particular entry are dependent on the class of object that the entry describes.

Evolving Security Management Standards

The DIB's entries are arranged in the form of a tree, known as the Directory Information Tree (DIT) in which the vertices represent the entries. Entries higher in the tree (nearer the root) often represent objects, such as countries or organizations, while entries lower in the tree represent people or application processes. The services defined in the X.500 recommendation operate only on a tree-structured DIT. Every entry in the DIB has a distinguished name that uniquely and unambiguously identifies it. Properties of the distinguished name are derived from the tree structure.

The Directory enforces a set of rules to ensure that the DIB remains well formed in the face of modifications over time. These rules, known as the Directory schema, prevent entries that have the wrong types of attributes for their object class, attribute values of the wrong form for the attribute type, and entries having subordinate entries of the wrong class.

The Directory also provides users with a well-defined set of access capabilities, known as the abstract service. This service provides a simple modification and retrieval capability, which can be expanded with local DUA functions to provide the capabilities required by the end users.

It is likely that the Directory will be distributed along both functional and organizational lines. Models of the Directory have been developed to provide a framework for the cooperation of various components to provide an integrated whole. The provision and consumption of the Directory services require that users (actually the DUAs) and the various functional components of the Directory cooperate with one another. In many cases this requires cooperation between application processes in different open systems, which in turn requires standardized application protocols to mediate the differences.

The Directory has been designed to support multiple applications, drawn from a wide range of possibilities. The nature of the applications supported will govern which objects are listed in the Directory, which users will access the information, and which kinds of access they will carry out. Applications may be very specific, such as the provision of distribution lists for electronic mail, or generic, such as the "interpersonal communications directory" application. The Directory provides the opportunity to exploit commonalities among the applications:

- A single object may be relevant to more than one application. Therefore, a number of object classes and attribute types are defined across a range of applications. Recommendations X.520 and X.521 contain these definitions.

X.200 Open Systems Interconnection—Basic Reference Model
X.208 Open Systems Interconnection—Specification of Abstract Syntax Notation One (ASN.1)
X.501 The Directory—Models
X.509 The Directory—Authentication Framework
X.511 The Directory—Abstract Service Definition
X.518 The Directory—Procedures for Distributed Operation
X.519 The Directory—Protocol Specifications
X.520 The Directory—Selected Attribute Types
X.521 The Directory—Selected Object Classes
X.219 Remote Operations: Model, Notation and Service Definition
X.229 Remote Operations: Protocol Specification

Table 1. OSI standards on Directory.

- Certain patterns of Directory usage will be common across the range of applications.

The December 1988 version of the X.500 incorporates changes that were made as a result of the Draft International Standard (DIS) ballot in ISO; the text that will actually appear in the CCITT Blue Book in late summer or fall 1989 may incorporate some changes, however. The 1989 Blue Books will be the definitive source. Table 1 depicts the family of newly defined X.500 standards as they appeared in the December 1988 document.

Directory Services

The present X.500 standards do not completely define all aspects of Directory service. Therefore, some capabilities are provided as local functions until a standardized solution is available. These capabilities include addition and deletion of arbitrary entries, thus allowing the creation of a distributed Directory; the management of access control granting or withdrawing permission for a user to access a particular piece of information; the management of the Directory schema; the management of knowledge information; and the replication of parts of the DIT.

Directory users (people and computer programs) can read or modify the information in the database, or parts of it, if they have permission to do so. Each user accesses the Directory through a Directory User Agent (DUA), which is considered an application-process, as shown in Figure 1. A DUA can reside on PCs, workstations, manufacturing devices, and other types of equipment. It may be a standalone program or a component of a larger program. The Directory provides services in response to requests from DUAs, which allow *interrogation* and *modification* of the Directory. The Directory always reports the outcome of each request.

Directory Interrogation Services include the following:

Evolving Security Management Standards

- **Read**, aimed at a particular entry, causes the return of all or some of that entry's attributes. Where only some attributes are returned, the DUA supplies the list of attribute types of interest.
- **Compare**, aimed at an entry's particular attribute, causes the Directory to check whether a supplied value matches a value of that attribute. For example, this can be used to carry out password checking, where the password held in the Directory might be inaccessible for read, but accessible for compare.
- **List** causes the Directory to return the list of immediate subordinates of a particular named entry in the DIT.
- **Search** causes the Directory to return information from all of the entries within a certain portion of the DIT that satisfies some filter. The information returned from each entry consists of some or all of the attributes of that entry, as with read.
- **Abandon**, applied to an outstanding interrogation request, informs the Directory that the originator is no longer interested in completing the request. Directory may cease processing the request and discard any results so far achieved.

Directory Modification Services include the following:

- **Add Entry** causes the addition of a new leaf entry (either an object entry or an alias entry) to the DIT.
- **Remove Entry** causes removal of a leaf entry from the DIT.
- **Modify Entry** causes the Directory to execute a sequence of changes to a particular entry. Either all or none of the changes are made, and the DIB is always left in a state consistent with the schema. Allowable changes include the addition, removal, or replacement of attributes or attribute values.

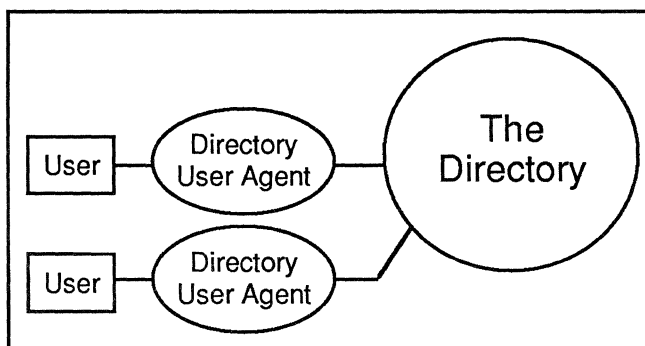


Figure 1. Access to the X.500 Directory.




Figure 2. X.500 function model of the Directory.

- **Modify Relative Distinguished Name (RDN)** causes modification of the relative distinguished name of a leaf entry (either an object entry or an alias entry) in the DIT by the nomination of different distinguished attributed values.

A Directory System Agent (DSA), an OSI application process that is part of the Directory, provides access to the DIB, DUAs, and/or other DSAs. DSAs use standard Directory Access Protocol (DAP) and, optionally, Directory System Protocol (DSP). A DSA may use information stored in its local database or interact with other DSAs to carry out requests. Alternatively, the DSA may request another DSA to help carry out the request. The DUA interacts with the Directory through one or more DSAs, although it is not bound to any one in particular. (See Figure 2.)

Authentication via the Directory

Many applications require objects participating in a session to offer some proof of their identity before they are permitted to carry out specific actions. The Directory also supports this authentication process. In the less complicated approach to authentication, called "simple authentication," the Directory holds a "User Password" attribute for any user that wants to authenticate itself. At the request of the service under consideration, the Directory confirms or denies that a particular value is actually the user's password. Because the Directory authenticates the password, a user needs only one, rather than a different password for every service. When exchange of passwords in a local environment using simple authentication is inappropriate, the Directory protects those passwords against replay or misuse through a one-way function. (See Table 2.)

The more complex approach, called "strong authentication," is based upon public key cryptography, where the Directory acts as a repository of users' public encryption keys, suitably protected against tampering. The steps that users can take to obtain each others' public keys from the Directory, and then to au-

Evolving Security Management Standards

Authentication Classes	Function
Simple authentication	Employs the user's unique name and password to provide a level of identification assurance
Strong authentication with secret key	Employs symmetric encryption to provide a level of assurance
Strong authentication with public key	Employs asymmetric encryption to provide a level of assurance

Table 2. X.500 Directory—three authentication classes.

authenticate with each other using them, are described in detail in Recommendation X.509. The United States government has at least three organizations that sponsor security standards: the U.S. Department of Defense (DOD); the Department of Commerce, Federal Information Processing Standards (FIPS); and the General Services Administration (GSA), which produces federal standards.

Department of Defense. The DOD sponsors both DOD and military standards. The Department's National Computer Security Center (NCSC) has produced many recommendations, some of which have become standards. Efforts such as the Secure Data Network System (SDNS), a National Security Agency program to develop standards for secure data networks, and Government Open Systems Interconnection Profile (GOSIP) should produce standards in the next few years. *[EDITOR'S NOTE: the X.500 Directory is not currently part of GOSIP.]* The SNS program integrates security features into OSI-based networks, using the concepts of OSI security discussed above. The program will issue standards that may then be used by vendors involved in NSA's Commercial COMSEC (Communications Security) Endorsement Program.

National Computer Security Center. In 1978, the National Bureau of Standards (now the National Institute of Standards and Technology) started to discuss the auditing and evaluation of computer security. In June 1981, the DOD Computer Security Center began operation. The success of the Center and the need to transfer security information to the commercial sector resulted in the expansion of the DOD Computer Security Center to the National Computer Security Center (NCSC). NCSC's mission is "to develop and promulgate uniform computer security criteria and standards." The NCSC has issued the publications listed below.

- Orange Book, a DOD standard for evaluating stand-alone computers for security, defines several differ-

ent levels for secure computers. In order of increasing security, these levels are D, C1, C2 (traditional) and B1, B2, B3, A1 (multilevel).

- Yellow Book is a guideline to the environments (external circumstances, conditions, and objects that affect the development, operation, and maintenance of a system) in which each Orange class offers adequate protection.
- Red Book is a standard for interpreting and applying the Orange Book requirement to the evaluation of networks.
- Green Book provides guidelines for the initialization, administration, and maintenance of password systems.

Federal Information Processing Standards. FIPS standards are published by the National Institute of Standards and Technology, which obtains input from other areas of the government and from the private sector. An example of a FIPS security standard is the Data Encryption Standard (DES).

Financial Industry Security Standards. Because a major concern is secure electronic transfer of money and securities, the financial community has devoted extensive resources to information security. The industry works through the American National Standards Institute (ANSI). The ANSI groups responsible for financial security are American Standard Committees X9 and X12. ASC X9 develops financial services standards; ASC X12 develops standards for business transactions. International responsibility for financial security standards is distributed between ISO Technical Committee 68 and Technical Committee 154.

The financial community is concerned with the following:

- Connection integrity (also called message integrity) is provided through a digital signature—an encryption method that guarantees that data has not been altered or destroyed. Financial standards use a message authentication code to provide this service. The ANSI X9.9-1986 and ISO 8730 standards define the process.
- For connection confidentiality, the financial community uses the Data Encryption Algorithm (DEA X3.92-1981), equivalent to the Data Encryption Standard, to protect data. Financial institutions use encryption to protect personal identification numbers (PINs) (ANSI X9.8-1982) and messages (ANSI X9.23).

Evolving Security Management Standards

- Access control management. An effort is under way in ANSI to develop a standard (ANSI X9.26-Draft) to protect computer sign-on information, such as passwords.
- Encipherment (Key) management. The financial community relies on a Key Management Center (KMC) to distribute keys to subscribers that must communicate. The center audits and assigns liability based on a subscriber's knowledge of the key. This is a necessary feature for the financial community. Communication among key management centers is under study so that subscribers of different centers can communicate in a secure manner.

Standards are fundamental to security and security management. A number of standardization efforts are near completion and should start to bear fruit within a few years.

BIBLIOGRAPHY

- A.L. Andreasson, *et al.*, "Information Security: An Overview", *AT&T Technical Journal*, May/June 1988, pp. 2-8.
- K. Barker, L.D. Nelson, "Security Standards—Government and Commercial," *AT&T Technical Journal*, May/June 1988, pp. 9-18.
- A.J. Bayle, "Security in an Open System Network: A Tutorial Survey," *Information Age*, Volume 10, No. 3, July 1988.
- W. Diffie, "The First Ten Years of Public-Key Cryptography," *Proceedings of the IEEE*, Vol. 76, No. 5, May 1988, pp. 560-577.
- Final Text of DIS 7498-2, "Information Processing Systems—OSI Reference Model, Part 2: Security Architecture," ISO/IEC JTC 1/SC 21 N2890, July 19, 1988.
- ISO/IEC JTC1/SC21, "Information Retrieval, Transfer, and Management for OSI: OSIRM Part 4—OSI Management Framework." Revision of DIS 7498-4, October 1988.
- Office of Technology Assessment, "Defending Secrets, Sharing Data," OTA-CIT-310, October 1987. □

Disaster Recovery Planning in an Integrated Network Environment

This report will help you to:

- Employ the different elements of communications disaster recovery.
 - Plan strategically and tactically for integrated network survival.
 - Survey disaster recovery resources and service providers.
-
-

In developing a business plan and strategy, planners must factor in the very considerable costs for disaster recovery in terms of both tangible and intangible loss. Planning for disaster recovery of facilities in an integrated environment presents a different challenge than what might have been the case just 10 years ago. Prior to divestiture, telecommunications disaster recovery was very much the province of the serving Bell Operating Company (BOC). In like manner, backing up a computer facility implied identifying an off-site storage facility for files and locating a similarly configured computer system. For many firms such solutions to disaster recovery are archaic and in themselves represent disaster.

Disaster recovery in an integrated network environment presents many new problems as well as opportunities for disaster avoidance as well as disaster recovery. In recognition of this fact, several vendors have established very comprehensive disaster recovery services. Some of these companies specialize in disaster recovery for computer facilities, while others specialize in recovering telecommunications facilities.

In today's operating environment, where companies operate multiple locations, each with computers, digital PBXs, remote access arrangements, and a variety of functional elements, all of which depend upon computer and telecommunications access, the need for proper disaster recovery planning is extremely

important. This report discusses the various elements that are involved and should be considered in the disaster recovery process. The model for this discussion focuses on a typical business situation involving multiple locations and an integrated network environment.

THE THREAT

In defense analysis, Department of Defense planners base their plans upon a perceived threat in order to identify the resources required to counter that threat. In today's operating environment, telecommunications facilities are the major arteries through which many corporations either deliver or conduct their business. For many companies, the loss of telecommunications and computer facilities can be measured in precise dollar amounts. For others, losing communications can only be measured in soft dollars, yet the effect can be significant. Therefore, planning for disaster recovery begins first with identifying and then estimating the impact that a particular disaster will impose upon a user's facilities.

A small formal wear clothing manufacturer operating an online order entry call management center was asked what a blackout of his telecommunications lines would mean. His reply was simple: \$70,000 per day in lost orders. A paper distributor was posed the same question and his reply was similar: about \$80,000 in lost orders. In both cases, those orders would likely go to competitors if the loss of communications continued for any length of time.

This report was developed for Datapro by Andres Llana, Jr. Mr. Llana attended Temple University, U.S. Army Signal School, and the U.S. Army Command and Staff College. He is a telecommunications consultant with Vermont Studies Group, Inc.

Disaster Recovery Planning in an Integrated Network Environment

Since 1987 there have been several major communications disasters, the most recent being a fire in the central switching facilities in Hinsdale, Illinois; prior to that, several major floods and a major earthquake in Mexico. The result of such events is always the same: major losses in business opportunities. For example, during the Mexico earthquake, the Bank of America and its affiliates in Mexico City estimated it lost 1 million dollars a day in float as well as penalty interest on cross-border loan payments and trade and foreign exchange losses. In May 1988, the fire that engulfed the Bell switching center in Hinsdale, Illinois destroyed fiber and digital circuits supporting over 50,000 lines. This disaster totally disrupted business for many large and small businesses operating through that center.

Obviously, major disasters can cause serious business interruptions. What are the levels of threat that can imperil a business and how can a business recover given such circumstances?

Line loss. In this scenario, line loss can result from the following:

- Major equipment component failure; i.e., line cards in a PBX, multiplexer, modem, or similar network terminating equipment.
- External physical damage to lines entering the premises.
- External damage to the local switching center.

Loss of power. In this scenario, power failure can be attributed to:

- Loss of internal equipment power supply.
- Loss of public power supply.
- Lightning strike.

Direct equipment failure. In this scenario, communications loss can be due to:

- Loss of a PBX.
- Loss of a multiplexer.
- Loss of a front-end processor.
- Loss of a mainframe computer.

Major facility loss. In this scenario, a major loss can be attributed to:

- A major flood.

- A major fire.
- Earthquake.
- Tornado, hurricane, etc.

Sabotage (human-made). In this scenario, either inside employee or former employee conducts some form of equipment or facility damage.

For all these disasters, managers can effect action plans to avoid or soften their impact. Disaster plans can be both short term and long term. For example, in the case of the earthquake in Mexico, Bank of America was caught completely off guard. However, a Bank of America fast reaction team used high frequency radio (short term) to reestablish voice communications while rebuilding land-based telecommunications facilities (long term). Bank of America has since instituted a number of disaster recovery plans that encompass both natural and human-made disasters. In addition, key Bank of America personnel have been trained as licensed radio operators to serve during a disaster, insuring prompt recovery of voice communications.

In the case of the Hinsdale fire, several companies were able to reestablish online computer facilities by moving to a remote cold site. All of these companies already had disaster recovery contracts in force with existing recovery centers; thus, prior planning paid off.

The most frequently reported sabotage in American business is attributed to disgruntled former employees. Therefore, disasters caused by sabotage can best be avoided through the diligent exercise of a well-documented security process designed to meet such threats.

THREAT ANALYSIS

In conducting the scenarios for defense, managers must take each threat individually to determine the resources required to counter the threat and the impact that it could make on the company's operations. For example, line loss due to a major component failure could last for an entire day. This type of threat is well within the means of the user to control, however, through a variety of self-help measures. External physical damage can result from flooding, cable severance, destruction of external CO distribution facilities, etc. These conditions are harder to deal with since they are outside the user's control. Again, there are some actions that users can take to offset the impact of such actions.

Disaster Recovery Planning in an Integrated Network Environment

Power failures are common occurrences and are handled routinely by the serving utility. However, the user must plan for such occurrences and adopt measures to counteract power outages. This might range from simple short-length UPS and battery backup schemes to power generators and sophisticated power regeneration sources.

Direct loss of specific major equipment components demands the rapid identification of off-site equipment and its relocation to the user's facilities, or the relocation of the user's personnel to cold sites, to bring replacement systems online for successful operations. In like manner, a major facility loss such as fire, flood, or other natural disaster requires in most cases the total reconstruction of the user's facilities. This type of threat requires greater personnel and equipment resources. Time to recovery is measured in weeks rather than days, which may apply in the previous scenarios.

Certainly, management must view the cost of disaster recovery as an insurance policy; however, care must be taken in developing an insurance premium that is cost consistent with the disaster's impact. Table 1 shows some common problems that could become disasters and remedial solutions.

STRATEGIC PLANNING FOR NETWORK SURVIVAL

Planning for network survivability requires a strategy for both equipment components and hardware components composing a total network. Many users assume that normal maintenance contracts supported

by their equipment providers cover all contingencies; this is not so! In fact, users should build provisions into their hardware acquisition and maintenance agreements for disaster recovery and system backup.

Preventive Measures

One of the most common problems in any computer or telephone operating environment is damage caused by water intrusion. Water falling on electrical equipment can cause immediate damage, shutting down equipment as fast as a lightning strike. All communications and computer facilities must be protected from water intrusion. Facilities at ground level should be elevated, have waterproof ceilings, and provide fast drainage to insure that any water intrusion can be quickly carried away. Facilities located on upper floors must use only waterproof ceiling fixtures to protect equipment from water falling from higher floors. Automatic water sprinklers should not be used for telephone and computer facilities. Such facilities should be equipped with Halon and fire extinguishers approved for use with electrical systems. In like manner, all wiring closets as well as the mainframe should also be protected from water intrusion. If there is any danger of uncontrollable water intrusion, water detection as well as protective coverings should be readily available in the facility room.

Users should not assume that facility air-conditioning and air circulation are adequate. Systems that continually overheat will degrade in performance and experience component failure. In some installations users have installed separate air handling equipment to support an isolated telephone room that cannot be supported by an existing facility. In all cases, if secondary air conditioners are deployed, they should be equipped with overflow pans to control condensation.

Fire protection should not be taken for granted. As a rule, most large-scale computer and telephone facilities are equipped with Halon systems. For smaller installations, however, dry fire extinguishers, e.g., Halon, CO₂, etc., should be available. Freestanding fire extinguishers should also be checked frequently. Large installations should have a "hot line" to the local fire department to arrange for inspections.

Power brownouts and outright power failure are common in many fast-developing areas. Lightning strikes are also a common source of equipment loss. It is not uncommon for a lightning strike to consume an entire PBX or some of its components, leaving the user without communications. Thus, every effort should be taken to install surge and lightning arresting equipment, all properly grounded, between the user's computer or telephone system and the power source. (See

Type	Impact	Approach
1. Line loss due to component failure	Partial loss in service	Maintain on-site major spares or backup component
2. Line loss due to construction intrusion	Total loss of lines 1 to 2 days	Off-site facility for partial contact, i.e., sales office, etc.
3. Loss of power due to lightning strike	Temporary loss of power	Maintain battery backup, UPS, power generator
4. Major equipment failure, i.e., PBX, computer, etc.	Loss of service	Contract for or maintain standby item of equipment
5. Major facility loss through fire or natural disaster	Loss of all facilities	Contract or maintain off-site cold or shell facility

Table 1. Potential disasters and remedial solutions.

Disaster Recovery Planning in an Integrated Network Environment

Report MT50-380-101, "Plant Protection and Hazard Management," for more information on grounding and lightning protection.)

Brownouts can also affect a PBX or key system, but in different ways depending upon the manufacturer. Managers should contact other users of their equipment to ascertain components that might be stored on-site as insurance against component failure. In some cases, faulty power can cause a PBX to lose its memory, making it necessary to restart the system, which in turn will produce long delays in telephone service. There are many vendors selling small-scale UPS systems that protect a system's memory and disk subsystem from sudden crashes caused by power failure. In addition, battery backup coupled with a small-scale UPS system can support a small PBX for four to eight hours. Larger PBX installations use much larger battery backup systems and can stay on-line up to eight hours. Depending on the location of facilities, some users find it necessary to utilize diesel generators as a means to back up their battery systems for both their computer and telephone systems, insuring the uninterrupted support of both.

With the integration of voice and data over the same wire pairs, the importance of internal cabling has increased. Major incursions into the cabling system integrity often produce major internal service disruptions due to the lack of sufficient documentation. For this reason, many users find that they are ill prepared to deal with a major disruption of their cabling system. To counter this problem, they must prepare and maintain accurate cabling documentation if they are to avoid service interruption caused by cable changes.

Small Network Strategies

A user must not believe that strategic network planning applies only to large networks. Network strategizing is just as important to small network users since, in many cases, they have more to lose in a network failure. Strategizing is mainly the backing up of critical functions that could not be supported in a time of service loss.

Small system users can find many ways to preserve the critical operations in their companies. For example, a total PBX failure can be avoided by having the main listed number and lines in the hunt group dual wired to special power failure units that will support one telephone for each of the associated lines. During any PBX outage, these specially located telephones will receive power from the central office and can be used to receive and make outside calls.

Some users will also terminate separate private lines at key locations around the company, which can be an effective strategy for several buildings in the same location. Each of these lines will terminate at an assigned location, bypassing the PBX and the principal wiring point of demarcation. Depending upon the central office facilities, some users will terminate Centrex lines instead of individual private lines, in effect providing a dual telephone network on the user's premises. The Centrex lines may be installed so that they bypass the Demarc, or they may be linked to the user's PBX via a tie line to the CO. In this situation the user would have a backup at some locations while having the effect of two separate telephone systems supporting the same facility.

Large Network Strategies

Large networks comprising several locations have a different set of problems than small networks, and managers must consider additional strategies to support them. In these networks it is not unusual to find that the data network portion has some form of backup while the voice network will have little if any backup support. Many data networks rely on dial backup circuits as a means of supporting a data network. Dial backup involves two dial-up modems, one at the host end and one at the remote end, as well as two telecom lines. When the primary private line fails, the user at the remote end uses the dial modem to reestablish direct contact with the front-end processor at the host end.

A major breakthrough developed in dial backup was the introduction of the V.32, 9600 bps dial-up modem. The more popular of the V.32 modems are those that are fully compatible with the V.22 and V.22 bis modem standards, which make it possible to use a mix of 1200, 2400, and 4800 bps modems at the remote ends. While this technique is effective for medium-sized networks, managers will find this approach much too costly for large networks.

Many networks that integrate voice and data over T1 links have a very serious backup problem since they risk all of their traffic on a single communications path. Backup strategies to support this scenario are many and diverse.

One popular T1 backup strategy is diverse routing, employing microwave, fiber, or similar media to bypass the serving utility. For example, some users in the Southwest have employed a special-purpose microwave-based carrier from whom they subcontract microwave bypass facilities. The carrier is unique in that its terminal locations are all located within a quarter mile of a Bell or AT&T primary

Disaster Recovery Planning in an Integrated Network Environment

servicing office. The users employ these bypass arrangements as alternate routes out of their facilities for both disaster recovery and traffic off-loading. In this way, the special microwave carrier network serves as a diverse route out of the user's local area, affording the user an opportunity to interconnect back into the public network at some other location. Since the carrier's terminals are all strategically located, they minimize the interconnection charges into the public or long-distance network. In this way, a user could suffer a complete failure of T1 links but still be assured of another route at T1 speeds.

Typically, the microwave link is a 192- or 288-channel, 2GHz or 6GHz digital radio over which the user can consolidate critical traffic. At the digital service center the user can also switch traffic to other OCCs like Sprint, MCI, etc. In addition, the user has the option in a disaster of switching all traffic to a remote hot site (see Figure 1). Note that the standby ACCUNET T1 link is configured to support access to the disaster recovery site.

Some users have been able to configure their networks to provide a measure of backup by providing multiple routing. Figure 2 shows a typical ring configuration consisting of four nodes with diverse routes "G" and "E" providing route redundancy. Additional network backup can be provided through a physical diverse route from one of the nodes into the public network. In addition, at least one node can be used to link the network to a remote hot site for standby recovery of the main computer facility. If the disaster recovery link is provided by either MCI or US Sprint, it is possible for the user to locate multiplexing equipment at the OCC's switching center to support a "drop and insert" application. This arrangement would provide for link sharing to support voice, data, fax, and video traffic as the need arises. Thus, the disaster recovery link, when not used for disaster recovery, can be used as an optional route for long-distance voice traffic as well as point-to-point data links.

Since the CO fire at the Hinsdale switching center, more carriers are sensitive to user needs for backup and recovery of their backbone links. Recently, some vendors have been offering both VSAT and small earth stations as premises equipment for the direct transmission of T1 traffic via satellite. In Figure 3, the user has installed VSAT terminals at each node, permitting it to back up the entire network in mirror image fashion. Typically, a user would interface this backup satellite link to the network through a T1 multiplexer. In some cases, depending upon the network design, it might be necessary to include echo cancellation equipment where integrating voice. In a disaster recovery scenario, channels can be activated

within 15 minutes and the user would pay only for the time. (See Disaster Recovery Resources and Service Providers at the end of this report.)

In some metropolitan areas, fiber companies are also providing limited bypass opportunities for users wishing to construct diverse routes out of their premises. Like the microwave strategy, fiber provides for convenient bypass of the vulnerable local loop into another switching environment. This diverse route, like the microwave link, can be linked into a long-distance carrier such as AT&T, MCI, RCI, etc., where the OCC switching center is accessible. Specialized fiber carriers are limited in their availability; therefore, they may not be a solution for some companies.

PBX and Network Equipment Backup

Voice systems (PBX and key) sales are not structured toward backup and recovery as are large-scale computer and minicomputer system sales. As a rule, it has been the common assumption that the voice system vendor would be a stocking distributor with ample spares on hand to meet any emergency. As the interconnect industry has matured, however, vendors are keeping fewer spares on hand than before. This situation lessens the likelihood of the vendor being responsive to a devastating loss. Of course, the large voice system manufacturers that distribute directly, such as AT&T, Northern Telecom, and Intecom, could support an emergency situation as part of their regular maintenance. As a matter of practice, however, they do not address such issues in their normal maintenance. Hence, users should press for disaster recovery support when negotiating a new system contract or renewing a maintenance contract.

Two firms have introduced disaster recovery voice systems on a limited basis. These are transportable, ruggedized containers housing complete systems for interfacing with CO and long-distance carriers as well as radio and satellite links. These units are modeled after similar systems that are part of the U.S. Military's telecom inventory. They are capable of being configured as Class V switching centers and can be easily airlifted into place or transported by all-terrain vehicles for rapid deployment.

Several vendors have reintroduced their microwave products as network restoration systems. These systems can be configured to support either analog or digital DS1 to DS3 and voice, data, or video carrier requirements. They are short-range, frequency agile systems in the 1.7GHz to 23GHz frequency range.

These systems are also container mounted and ruggedized for fast deployment, borrowed again from the

Disaster Recovery Planning in an Integrated Network Environment

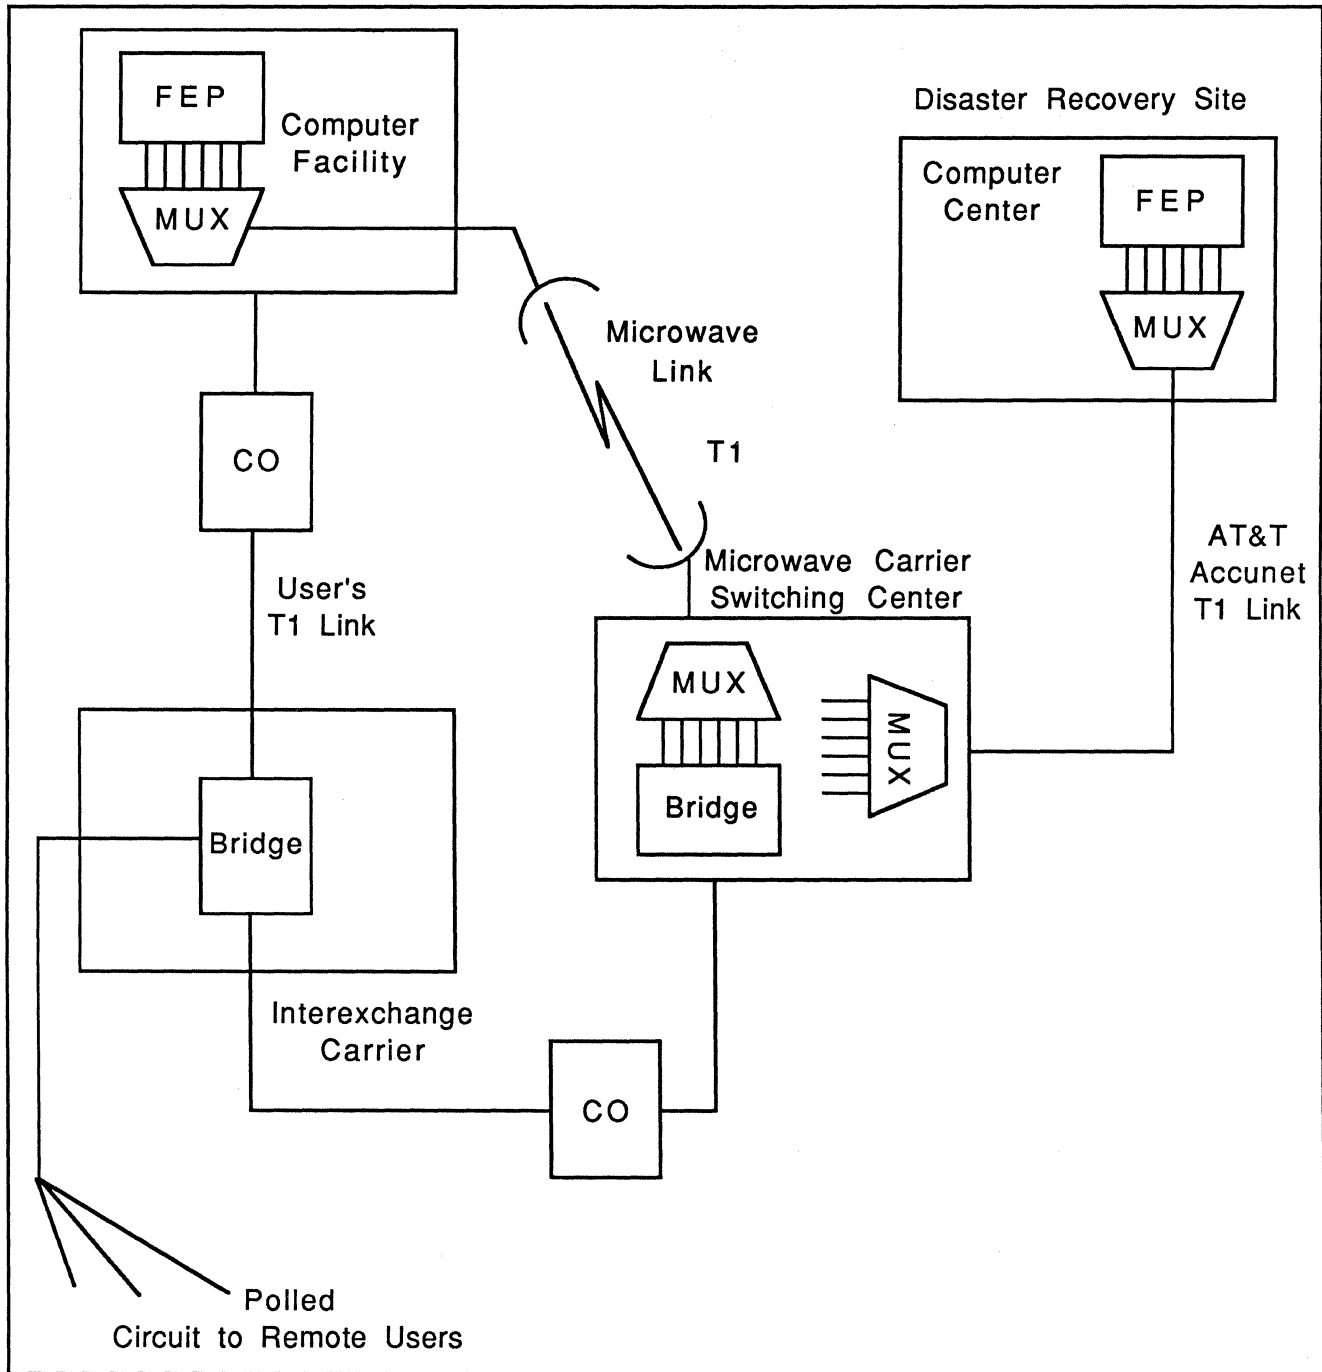


Figure 1. Microwave bypass link.

military, and focus on the rapid reorganization of telecommunications networks in a state of emergency. This self-help concept is not new to military planners; however, it is a relatively new concept for civilian planners who need to maintain vital telecommunications links.

COMPUTER FACILITY RECOVERY

Mainframe computers, front-end processors, and network terminating facilities such as matrix switches, modems, and multiplexers must be accounted for in a disaster recovery plan. Network terminating equipment such as T1 multiplexers are generally configured with full redundancy. This strategy, plus the maintenance of key multiplexer spares, will insure that these

Disaster Recovery Planning in an Integrated Network Environment

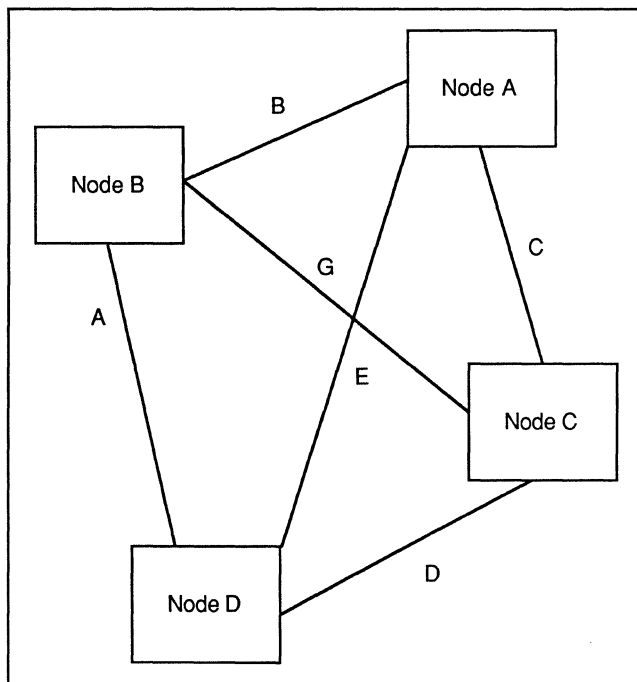


Figure 2. Network with diverse routing.

units maintain 99 percent uptime. In addition, some computer facilities maintain spare modems for substitution where specific units fail. A failed T1 multiplexer, however, can bring down an entire network segment. In this regard, supercritical network nodes should be backed up by "hot" standby multiplexers, assuring users that any choke point in their network is backed up.

Depending upon a network's configuration and the level of failure, a disaster will not be declared until the full potential of the user's mainframe computer facility has been compromised. There are several scenarios that may be followed in such a circumstance, but the philosophy for system backup has changed with the advent of data communications. In years past, backup meant moving tape and card files to a new computer center where data processing continued. While this technique may be followed where a data center has been totally destroyed, present-day data communications techniques streamline the process. Depending upon the applications, large users may revert to a "hot standby" site where all remote terminal locations are automatically switched without disrupting the end users. In this scenario, a user may have more than one data center integrated into a fully redundant network in terms of network routing and computer system configuration. The hot site in this network configuration would simply be one of several computer centers that might serve as a development

center during normal operations but configured to support any of the applications areas common to network users' needs.

In some circumstances federal government agencies, large corporate organizations, and financial institutions will use this technique to support national information networks. As a matter of practice, switching between these facilities will occur on a scheduled basis to allow for system maintenance. This can be viewed as disaster avoidance, since network downtime might only be experienced by users at the network extremities.

Companies that do not have multiple computer facilities must rely upon service centers specializing in disaster recovery services or cooperative agreements between users of like computer facilities. These "hot site" recovery centers offer users an array of equipment which is always available 24 hours per day on a subscription basis. In cases of major disasters such as a fire or flood, a user may declare a disaster and begin operations at the recovery center. Depending upon the user's configuration, recovery to normal operations at a hot site may take from one to several days depending upon the size of the operation. The recovery vendor site will be supported with extensive voice and data communications facilities which can be integrated into the user's own network.

Some users may also elect to maintain a "cold" standby site with a disaster recovery vendor. A cold site can be a facility co-located at a user's facility or at the vendor's location. The cold site will be fully equipped except that the systems located there may not be fully operational and require initialization of services. Usually, a vendor-owned cold site will have an inventory of computer equipment such that a user can configure a system exactly as the one that has been lost.

In less critical operations, such as manufacturing, a cold site can be used where hardware facilities may be on standby and operations can be implemented under less critical time constraints. In other circumstances, where the user must plan on a longer term "temporary" relocation, a shell site may be retained on long-term lease. All support facilities, such as electrical, telecommunications access, and environmental controls, will be in place. The user will then lease equipment on some short-term basis to support operations pending construction of a new permanent facility. Some vendors offer mobile cold sites, which are mobile units combining both computer facilities and people space. These take a bit more time to relocate but are configured to exactly the customer's needs. Mobile cold sites are moved to the customer's loca-

Disaster Recovery Planning in an Integrated Network Environment

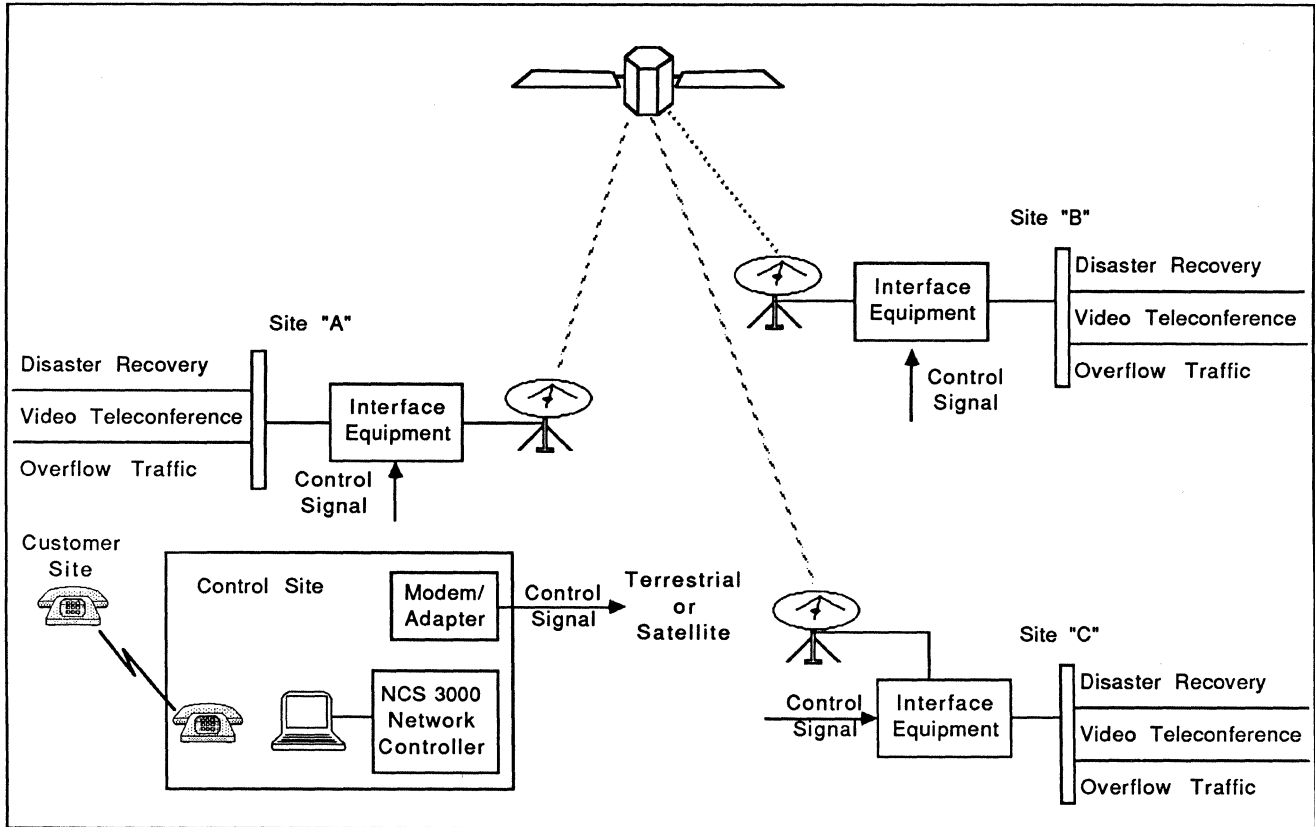


Figure 3. VSAT disaster recovery links.

tion of choice and then interconnected to both its internal and external network as appropriate.

In concert with this planning, the user will also maintain off-site secure storage for mass storage files and tapes that reflect current operations. During a disaster, these files would move from secure storage to the disaster recovery site.

SPECIALIZED OPERATIONS

Due in part to new banking laws, specialized centers have been set up on both a lease and shared basis to support check clearing operations. These centers may be existing operational sites that are brokered by a recovery vendor or covered by mutual agreements. Recovery vendor hot sites are equipped with check processing/scanning equipment. These systems can be configured to replicate the client bank's systems to insure prompt and accurate check processing in time of disaster. Recovery center telecom facilities allow the user to communicate directly with the corresponding Federal Reserve Processing Center.

Call management systems are unique since many of these are now supported by voice response systems and incoming 800 WATS-type lines that can be re-

routed automatically or on demand. These centers support online telemarketing operations, customer service centers, and similar operations. Basic call management system components are Automatic Call Distribution (ACD) systems, Bulk Incoming Call Facilities (800 lines), Voice Response Systems, and mainframe access. Backing up such operations to avoid a disaster can be accomplished in several ways, first by using the automatic call routing features of AT&T's 800 service to route calls to another service center. A second strategy involves rerouting calls to a contracted telemanagement center that will then be staffed to support the client's service requirements; however, contracted telemanagement centers may not have sufficient connectivity to access a user's remote computing facility. For example, many telemanagement operations use asynchronous computer systems in their normal operations. For this reason, special interface requirements may have to be established, such as protocol converters, in order to maintain a seamless interface during a crossover operation. Where access to a voice response system is a requirement, special arrangements may have to be made for a relocatable system. Voice response systems will usually support standalone databases which may or may not be up/downloaded on some scheduled basis, and could prove to be a major stumbling block for a seamless crossover of a call management center.

Disaster Recovery Planning in an Integrated Network Environment

DISASTER RECOVERY PLAN OUTLINE

Overview

Developing a disaster recovery plan is a detailed process that, once accomplished, must be reviewed and updated as often as the company adds new products, new facilities, or new technology or undergoes major organizational changes. A disaster recovery plan is like a blueprint for a house. The plan must, within the scope of the capability to be reestablished, be sufficiently complete that the lowest operating echelon could move to an entirely new facility and begin operations on a "business as usual" scenario. The plan follows the WHO, WHAT, WHEN, and HOW of any instructional manual: Who is involved, What resources are involved, When are these assets deployed, and How are they to be deployed. The plan must basically outline people issues, equipment resources, and detailed operating instructions; nothing can be assumed. It is the company's *tactical* battle plan for a "retrograde movement" that will facilitate the relocation of assets.

The Recovery Plan

Criteria for Disaster

Establishing criteria for a disaster is based upon the cost to the company of being deprived of that resource for any period of time. Thresholds must be established for each operation within the company. Restoration of service thresholds must be a part of this planning. For example, can the functional operation be without telecommunications or computer support for an hour, a day, or a week? What would be the consequences of not supporting a function for a period of time? How many other functions within the company would ultimately be affected by a breakdown in services? All of these issues and others unique to the business' critical needs must be addressed and prioritized when developing the criteria for declaring a disaster.

Notification

Who within the organization is responsible for declaring a disaster, and what steps are to be taken to put the plan into effect? It is important to establish disaster verification methods within the notification process if the disaster threshold might be backed by some form of insurance. Associated with the notification process are the contact names and telephone numbers of persons to be notified in declaring a disaster. As a

security measure, there should be several levels of notification to avoid false starts. The notification process is also a form of control in which a series of contacts must be made that lead to the final contact with the disaster recovery support vendors.

Escalation Plan. An escalation plan will be part of the notification plan. It includes all carrier contacts and their management. Every attempt should be made to obtain and include test center numbers down to the test console desk level. Vendor maintenance support personnel for each element of network hardware should also be included with up to three levels of management and above. In addition, the user's public relations resources should not be overlooked in the disaster recovery process. The public relations group can serve as an asset to effect pressure on an uncooperative vendor critical to recovery. All information contained in this portion of the plan must be reviewed and checked for accuracy at least quarterly since organizations are subject to change.

Disaster Recovery Team Organization

Organizing the disaster recovery team should be limited to just those hands-on personnel required to put a selected function into play. One central player should have the authority to act on the company's behalf at all times with the full backing of corporate officers. This is important where disaster recovery operations are conducted on a 24-hour basis at an off-site location. Extraneous personnel should be left out of play as they may only serve to confuse the issue. The plan then must incorporate all of the names and telephone numbers of both principal and backup members of the recovery support team as well as their vendor counterparts. There should always be primary and alternate players on a disaster recovery team; i.e., Red and Blue team. These member lists must be supported with asset lists outlining all of the assets required to support the team when deployed. No assumptions should be made, and planning should incorporate the level of detail to insure the team's self-sufficiency during the recovery process.

If equipment is required to support the team, the prepositioned location of all equipment and related tools must be documented and made part of a control list. For example, if a team member's responsibility is to initialize NCP on a front-end processor, then his/her equipment list would include the following: location of a console, portable terminal, copy of the current version release being used, current operating instructions, written instructions for a step-by-step initialization procedure (must be tested at least three times per year or as often as system changes are made), as well as any other support tool needed to

Disaster Recovery Planning in an Integrated Network Environment

insure successful initialization of the current NCP version. This type of detail must be covered for every operation and will insure a smooth system recovery.

People Support. No assumptions should be made as to employee availability. Alert lists must be compiled so that key personnel can be contacted during a disaster. Team members must know where to contact at least one other team member at all times. Alert lists can be tested periodically by having practice alerts that test the validity of telephone numbers and the availability of transportation to a predetermined assembly or departure area. Practice alerts should test the travel time to the recovery site to reveal problems in notification, transportation, and backup availability. All details concerning travel and transportation to the recovery site(s) must be accounted for in the plan. Advanced arrangements must be made for air transportation, land travel, housing, and meals. Relief arrangements should also be part of the plan so that personnel can be relieved after long periods of duty. Backup team plans must also be part of the recovery plan to insure continuity of all operations, particularly where recovery will involve long durations.

Disaster Recovery Hardware Plan

Every disaster recovery plan supporting an integrated voice and data network will require a hardware support plan. This includes test equipment, replacement spares, whole end items such as modems and multiplexers, backup PBXs, controllers, consoles, PCs, etc. The prepositioned location of all hardware must be identified with detailed instructions for locating all prepositioned locations. For example, if a hot site is being used as a mainframe computer recovery center, then the exact recovery center layout must be part of the hardware plan. Included would be all relevant location information details down to the location of 110 AC electrical outlets. If a vendor is under contract to provide an off-site PBX, multiplexer, or spares, the vendor's equipment location must be known to the user so that user personnel can be dispatched if it becomes apparent during an emergency that the vendor will not be able to deliver in a timely manner.

Recovery Site Plan

The recovery site plan must detail the recovery team's assembly and their required support at the recovery site. It is a detailed list of activities that must be followed in sequence in order to reestablish computer support in the minimum amount of time without wasted effort. This information is usually contained in a large loose-leaf book complete with a

series of appendixes representing additional operating instructions. Movement instructions concerning the removal and relocation of all off-site data must be identified, together with the persons or vendor responsible for the relocation operation. Procedures must be in place to be certain that all data can be validated and is properly identified.

An accurate and current site layout must be part of the recovery plan. The site layout should show the location of all equipment to include the CPU, FEP, DASD, matrix switch, multiplexers, modems, support CRTs, etc.

If matrix switches, dial backup modems, and other communications equipment are integral parts of the recovery process, then no assumptions should be made as to their location availability at time of need. For example, the required configuration for the matrix switch should be reflected in both written and magnetic media (i.e., disk). The multiplexer configuration should also be documented to be certain that all channel assignments match to the matrix switch. If dial backup modems are to be used, those units associated with the front-end processor must be identified. Premarked labels bearing the company's name and the associated device should be prepared for pre-identification of all equipment to be used at the recovery center at the time of recovery operations. Operating instructions should be part of the site plan. No assumption should be made as to the availability of experienced operators. Instructions should be included for the operation of any dedicated CSU equipment, multiplexers, modems, matrix switches, peripheral devices, and DASD and CPU consoles.

If communications lines must be rehomed at the user nodes, then detailed instructions for rehoming should be part of this plan even though they may be part of a communications vendor's support requirements. The remapping of multiplexers and other related equipment at the user's nodes should also be included as an appendix item. This also would include remote coordination plans and details for insuring that remote controllers are brought online via dial-up modems if these devices are to be used. Personnel at remote sites must be trained to bring up remote controllers using the dial-up modems. Procedures should also be in place to pretest any T1 links that may play a critical role in supporting the user's remote locations through the recovery facility.

The site plan should be considered a play in which each team member has a definite script. There is no room in time of a disaster for "ad libbing"; a team player must know where to go and what must be accomplished once the team has been assembled at the recovery center.

Disaster Recovery Planning in an Integrated Network Environment

Dress Rehearsals

A disaster recovery plan becomes worthless if the user attempts to exercise the plan and finds that it does not work. Therefore, part of the disaster recovery process should include frequent alerts with full system recovery testing. All of the disaster recovery vendors offer disaster recovery test periods as part of their service. During these tests, the user invokes the disaster recovery plan to determine effectiveness if it becomes necessary to restore services off-site. Some users find that their plans cannot be put into full effect for a number of reasons; however, it is important to check all of your plan's subcomponents on some regular basis. For example, practice alerts can be held to determine if alert lists are accurate. Scheduled tests can be run at the recovery site to time out the accuracy and recoverability of the front-end processor initialization instructions. At other scheduled tests, the instructions for validating and loading the data base can be checked to determine the accuracy of these instructions.

Some users run a practice session in which their entire system is brought online but no traffic is put on the network. At least once a year the entire user network should be brought up for one period during which all of the significant applications to be supported in a disaster mode can be tested. This process will identify network problems that may not show up until a true disaster occurs. For example, unanticipated application design problems may develop when the network is reconfigured during a disaster, and telecommunications resources may be limited. Furthermore, certain remote locations may not be supportable under certain conditions which may not show up until a time of disaster. Therefore, to minimize problems occurring during a recovery period, every effort should be made to test out the entire plan in as real a scenario as possible.

SUMMARY

For many years, users have taken for granted the ongoing operation of their computer and telecommunications systems. Users have become complacent, assuming that telephone services will always be available. Very few plan for their own mortality; however, it is important that you learn to think "tactically" if your business and that which supports your livelihood is to survive. Each year we are confronted with some form of potential disaster. If your planning is complete, you will be able to survive to some level of recovery. At present, off-the-shelf recovery services are limited and much depends upon the user's resourcefulness. In this report we outline some of the basic elements related to the disaster recovery plan-

ning process. In addition, we identify some of the off-the-shelf vendor support available. Users should not fail to recognize their vulnerability by not planning for survival; disaster recovery planning is a vital part of the business plan.

DISASTER RECOVERY RESOURCES AND SERVICE PROVIDERS

Network Disaster Recovery Resources

One resource that has been around for some time has been AT&T ACCUNET reserve T1 service. This service provides a user with a standby T1 circuit as a direct access link to a disaster recovery center. Commercial recovery service centers such as Comdisco, Sungard, Hotsite, and Corporate Computer Services, which are discussed later, have incoming ACCUNET T1 reserve service available to their clients. AT&T does not offer a disaster recovery service for its voice systems; however, its maintenance policy supports rapid deployment of any AT&T system that may be destroyed as the result of a disaster.

Recently, disaster recovery communications programs have been announced that may be worth considering. These are satellite (VSAT) services that provide for on-demand or interactive services to support point-to-point transmission at the T1 or 56K bps level. Contel ASC provides this service via its Channel Management Services. Channels can be established once the VSAT terminal (satellite earth station) has been installed; activation takes about 15 minutes. Transmission applications include voice, data, fax, and video; however, echo cancellation equipment may be required to support these applications. Contel ASC provides a turnkey consultation support service that includes all engineering and design aspects. Monthly lease for the earth station is \$2,500 per location plus \$300 per hour for airtime. There is also a network control fee of \$1,000 per month covering the user's network; custom pricing is available to support special requirements. The contact for this service is Luis Valencia (301) 251-8479.

CertainT-1 is GTE Spacenet's disaster recovery product. This is also a VSAT-type program offering both interactive and dedicated point-to-point service on demand. CertainT-1 also provides a dedicated 56K bps channel over the same earth station that can be used when not in a disaster mode. The CertainT-1 package is preengineered and includes echo cancellation equipment. CertainT-1 provides for both disaster recovery and scheduled service. As part of its package, GTE provides for quarterly testing of the installed system to insure readiness in case of a disaster.

Disaster Recovery Planning in an Integrated Network Environment

Monthly costs range from \$2,000 to \$2,500 for the VSAT earth station plus \$250 per hour for scheduled time with a \$1,000 per-day ceiling. Disaster recovery time costs \$750 per hour with a \$4,000 per-day ceiling. The contact at GTE Spacenet is Mary Henry (703) 848-1522.

Loral Terracom has introduced a packaged microwave system called InstaCom for DS1 through DS3 transmission. It can support full-duplex, point-to-point transmission in support of voice, data, or video applications. There are several available versions of the system which can be customized to the user's needs, including fixed plant-type configurations as well as containerized self-sufficient modules which can be rapidly deployed. Cost ranges from \$80,000 to over \$200,000 for a full-duplex system that is field deployable. Several systems can be used back-to-back to establish a relay system to extend the microwave signal. As part of the user's disaster recovery plan, the user must file with the FCC for a frequency clearance. In addition, the user would need at least two test transmissions per year to retain the frequency clearance for use during a disaster recovery. The contact for Loral Terracom is the Marketing Department at (619) 278-4100.

There are two companies offering specific programs aimed at disaster recovery of voice communications. These companies specialize in supporting the local serving utilities and government agencies in disaster recovery. One is Rotelecom, an affiliate of the Rochester Telephone Company, which provides an engineered, tailored disaster recovery support system for specific user requirements. Rotelecom's engineered service is designed around a mobile support system for both voice and data communications as well as radio communications. A fully ruggedized system, it is capable of rapid deployment in all kinds of environments. At present, it supports governmental agencies under special contract. Pricing is based upon the user's exact requirements. A contact for this program is Gene Kollmeir (716) 274-5480.

Redcom, located in upstate New York, manufactures test products for the telephone industry. It recently introduced a modular stacked unit that can support up to 48 ports. The system can be carded as either an analog or digital voice system. A single unit measures 26 by 18 by 32 inches and weighs approximately 100 pounds. The unit (SBX 384) can be easily located on a desk or table with administration handled through an RS-232-C port and any asynchronous terminal or any DTMF telephone. It supports all types of trunks, e.g., CO, 4-wire E&M, TIE, or FX; however, it will only support 2500/500-type straight sets. The basic system can be powered by a small power generator, or it will work off of 115/230 V AC where available. At

idle, the unit draws less than 75 watts from the available power source. These units can be purchased with a special modular container equipped with a complete support system. The basic systems can be stacked and interconnected to provide sufficient port capacity for a Class V office. These units have been purchased by some Bell Operating Companies and the federal government, including DOD and the Federal Emergency Management Agency (FEMA) for prepositioned telecommunications support to insure rapid deployment in time of emergency. Individual 48-port units sell for approximately \$23,000 to \$25,000 (unit price quantity one). The contact at Redcom is Dave Ross (716) 924-7550.

Disaster Recovery Service Providers

Comdisco

Comdisco is one of the largest and most comprehensive service providers. It offers hot, cold, and shell locations in seven geographic areas throughout the U.S. These include Wood Dale, Illinois; Atlanta, Georgia; San Ramon, California; Cypress, California; Grand Prairie, Texas; Carlstadt, New Jersey; North Bergen, New Jersey; Cranford, New Jersey; and Bridgeport, New Jersey. Comdisco provides support for IBM 4300 and IBM 33XX CPUs, 3705, 3725, 3745 FEPs, plus some Digital Equipment Corporation and Tandem computer equipment. It also remarkets check processing centers offering IBM 3890 MICR reader/sorters for check processing in Colorado, Illinois, South Carolina, California, and New Jersey. There are mobile 327X cluster controllers for deployment to customer remote sites to tie them into the service center. It also maintains vault service in Carlstadt, New Jersey. Comdisco provides cold sites and shell sites and will lease systems configured to its customers' requirements to support these leased shell sites. The centers all have sophisticated voice and data communications arrangements with VSAT backup. In addition, Comdisco offers its Comline2+, Muxlink, and AT&T ACCUNET communications network arrangements to support access to its service centers. Comdisco also provides relocatable shells which are preengineered to a user's requirements. These "shells" are freestanding, panelized buildings that can be transported to a user's site, complete with all environmental controls and utilities in place.

The cost of this service depends upon square footage requirements and is based upon the exact user requirements. Comdisco prices range from \$1,000 to \$20,000 per month for service access. In addition, there are notification and daily use fees as well as other service fees based upon the type of access and

Disaster Recovery Planning in an Integrated Network Environment

use of other facilities. Digital Equipment and Tandem system fees range from \$1,800 to \$4,000 per month plus notification and daily use fees. The Comdisco number in Rosemont, Illinois is (312) 698-3000.

Hotsite

Hotsite, Inc. maintains sites in Niles, Ohio; Raleigh, North Carolina; and Boston, Massachusetts. Its service is similar to that of Comdisco and caters to the midrange computer user. It supports the IBM 30XX and 4300 CPU as well as IBM 3705, 25, and 45 FEPs. It also has both 3890 and 1441 MICR check/sorter support and offers both cold and shell sites that are computer ready. Unlike Comdisco, it does not lease equipment but will act on a customer's behalf to acquire equipment for a shell site. Hotsite provides an innovative, mobile, relocatable shell support program aimed at the small user market that employs Digital, Wang, Data General, and other similar mini-computer systems. Hotsite will supply any computer configuration common to users' needs. Users contract for this service over five years and, in the event of disaster, a fully configured system, including the required people space, is delivered to the customer site. There is a full range of telecommunications support as well as backup telecom facilities for integration with the user's network. Pricing ranges from \$900 to \$2,500 per month for service access plus disaster activation fees that range from \$1,000 to \$5,000 plus daily use fees from \$250 to \$500. The relocatable shell service costs about \$495 per month and is de-

pendent upon space and computer configuration needs. Activation and daily use fees also apply. Contact is through the marketing department at (216) 652-9624.

Corporate Computer Service

Corporate Computer Service in New Hudson, Michigan offers a hot site at New Hudson supporting IBM 4300 and 33XX CPUs as well as 3705, 25, and 45 FEPs. There is a full range of telecom access available. Contact is through the marketing department at (313) 486-2110.

SunGard

SunGard Recovery Service operates a MEGA hot site in Philadelphia, Pennsylvania. This site supports IBM 30XX and 4300 CPUs as well as 3705, 25, and 45 FEPs and Comten 3595 FEPs. The FEPs can be front ended with a matrix switch. User-supplied multiplexers can be installed in the center to support direct user access from remote locations to recovery center systems. The center is supported with a sophisticated telecommunications access arrangement that includes fiber, ACCUNET T1 facilities, and backup telecom support. There is a mobile System 38 that is also available. Contact is through the marketing department at (215) 676-0600. □

Problem Management: Using the Help Desk

This report will help you to:

- Establish a Help Desk for your own network.
 - Determine basic Help Desk functions.
 - Design a structure for network problem records.
-
-

The goal of network management, regardless of network size or equipment, is to keep the network up and running smoothly. As the size of a network increases, so do the problems. In a network with under 1,000 devices, about 5 percent of the end users will call for help every day. This figure increases to 20 or 30 percent in the largest networks. While not all problems result in lost work time and, therefore, lost revenue, many may. Other difficulties may only annoy end users and increase their frustration with the network. The effect of frustration is hard to measure in dollars, but its cost in lost productivity is real. See Figure 1 for an estimate of calls per day related to the number of devices.

Keeping a network cost effective and efficient, as well as reducing user frustration, requires quick resolution of network problems. One solution to effective problem management is a Help Desk. Conceptually, a Help Desk is a clearinghouse for reporting and resolving problems. Each time a user calls, the Help Desk tracks the difficulties—documenting what they are, how they are resolved, whether they occurred before, and any other pertinent information. Problems can take many forms: hardware, software, documentation, procedures, security, etc.

This report was developed exclusively for Datapro by Judy Van Tijn, a free-lance writer and systems training consultant. Ms. Van Tijn has eight years in the PC and data processing industry, with experience in editing, developing documentation, and creating PC training courses.

The Help Desk can be as simple as one person answering calls or an electronic mailbox for messages about problems. It can be as complex as the network requires.

HELP DESK FUNCTIONS

A Help Desk's exact functions must be determined at the beginning of network planning. Depending on the network, these functions can range from simply resolving problems to overseeing the entire network, including inventory and maintenance procedures. Without procedures for reporting, recording, and tracking problems, however, valuable time can be lost in redundant effort. Two or three people can be chasing the same problem and not know it. The problem could have occurred before and been solved, and the information later lost. Table 1 lists some common network problem categories and the job titles of staff members who typically receive the calls, in the absence of a central Help Desk.

Centralized reporting and recording of problems leads to better corporate management of networks. According to Mr. Gypsy Munoz, data control administrator at Affiliated Banks Service Company in Colorado, a Help Desk combined with a problem tracking system definitely provides the best solution. In any organization, problem management is key to quality assurance. Particularly when network users are also customers, the system loses integrity if problems are not resolved quickly and efficiently. A good problem management software package can shorten response time to five minutes or less by automatically alerting

Problem Management: Using the Help Desk

Devices	Problems/day	Cost/year
100	5	\$62,500
500	25	\$312,500
1,000	50	\$625,000
5,000	250	\$3,125,000
10,000	880	\$11,000,000
50,000	11,500	\$143,750,000
100,000	40,000	\$500,000,000

Source: Peregrine Systems

Figure 1. This figure shows the average costs associated with outages due to network problems. The cost calculation assumes that an average call translates into an outage of one half user-hour, that it takes one half technician-hour to resolve, and that a full-time equivalent person (FTE) costs \$100,000 per year.

the right person to handle a problem and keeping a database of past problems and resolutions. Furthermore, some packages allow all aspects of network inventory and maintenance to be correlated through the Help Desk. This can assist in identifying recurring problems and repeated failures and aid preventive maintenance. Various software packages are available to automate Help Desk functions, but network managers can use any relational database for simply tracking problems if it is carefully designed.

ESTABLISHING A HELP DESK

Key elements in establishing a Help Desk include:

- A clear understanding of Help Desk functions
- Management support
- A well-designed tracking method

As discussed previously, basic Help Desk functions include centralized reporting, standardized recording, and comprehensive tracking of trouble calls. Other more sophisticated functions, such as controlling inventory and maintenance procedures, may be added.

When seeking management support, it is important to stress that Help Desk setup costs are offset by savings in network availability.

A well-designed tracking method should be based on the problem life cycle and reflect specific Help Desk functions. A problem record is opened when someone calls into the Help Desk asking for assistance. The record is updated to maintain an audit trail of everything associated with the problem and, when the problem is resolved, the record is closed. Information about that problem is then available for reporting, analysis, and historical reference. The quality of in-

formation gathered during this process will affect the Help Desk's ultimate effectiveness.

It is critical to record all problems, no matter how trivial or how quickly solved. Program features that minimize keystrokes and make entry quick and simple maximize the likelihood that problems will be entered. Unrecorded problems can hinder network management and reduce the Help Desk's effectiveness. If all calls are not logged in, important information can be lost, and easily resolvable and preventable problems can wind up generating the majority of calls. For example, one organization may allow users to change passwords every week, creating confusion among end users and swamping the Help Desk with calls. Restricting password changes to once a month could provide adequate system security while reducing the number of problem calls.

BASIC CONCEPTS

The Problem Record

A problem record should contain fundamental information about the problem: what it is, its category, to whom it is assigned, when it was reported, the configuration of affected equipment, who reported it, how it was resolved, any previous problems with the item, and a unique identifier for labeling the specific problem. Determining exact fields for any given problem template depends on the specific network being implemented and the functions it will perform. If inven-

Problem	Person Receiving the Problem Call
On-line ABENDS	Operators
Batch ABENDS	Applications Programming
Lost reports	I/O or Production Control
VTAM vary ACTIVE	Network Control
TSO cancel userid	Operators
Bugs (system)	Systems Programming
Bugs (application)	Applications Programming
B37 ABENDS	DASD Management
Dataset lockouts	Operators
Terminal/line problems	Network Control
Forgotten password	Security
Tape problems	Tape Librarian
Scheduling problems	Operators or Production Control
PC problems	Infocenter
End user help	Customer Support

Table 1. Some typical categories of trouble calls. MIS managers in medium to large networks are rarely aware of the actual number of trouble calls they receive. This lack of awareness is because the calls are not logged and are not received by a central Help Desk.

Problem Management: Using the Help Desk

tory, configuration, or change information is part of the same database, fields to track this information must also be defined.

Whenever possible, problem information should be automatically filled in from the database. For example, when a person enters a problem about a terminal, the program should automatically fill in additional information about who owns the terminal, what it is supposed to do, what line it is on, what software it is running, and any previous problems. Automatic data entry accomplishes three things: it guarantees data integrity, minimizes keystrokes, and creates/stores a "snapshot" of the network at a particular moment in time.

Typically, a problem is reported about a device that may then be moved to another system or line. In such a situation, the assignment group cannot fix the device because it has been moved. Alternatively, a device fault may be secondary to another problem already reported; for instance, it may be connected to a controller that is down. With accurate data automatically filled in, the problem record preserves the network's status at the time the problem occurred.

Network personnel should be able to search online any field in the problem template. If it were possible to know in advance what problems would occur, they would not become problems. Furthermore, it is not possible to predict criteria/fields on which people will want to search. The greater the flexibility, the more information available in an emergency. Although not necessarily efficient, many answers can be found by thorough searching.

Problem Categories

Every network problem must be categorized. A typical Help Desk employs 10 to 30 categories for defining problems, which reflect specific network elements.

Each category has its own data entry template, a set of fields that ask for specific information pertaining to that category. Different problem categories have different severity levels. If the mainframe or centralized server is down—critical problems—different people should be notified than if the problem is a single terminal failure. Different categories, therefore, imply different procedures.

It is often necessary to change a problem's category after one has been assigned. It may have been reported as a terminal or printer problem but actually be a line or software problem. It may also be necessary to reopen a problem that had been resolved.

Both these capabilities should be designed into the problem management system.

Assignment Groups

Assignment groups are responsible for fixing certain categories of problems. Programmers, hardware specialists, and communications maintenance all constitute different assignment groups. Different kinds of problems are handled by different specialists. As soon as a problem record is opened, the relevant assignment group should be notified. Ideally, this happens automatically as soon as the problem's category is entered.

Alerts

The Help Desk system should track a problem through alerts, of which there are at least three kinds:

1. An escalation or activity alert. This occurs when a preset amount of time passes and an opened problem remains unresolved. The amount of time is determined by severity and other parameters. The assignment group is notified that the problem is overdue. These alerts escalate at preset intervals with higher and higher levels of management being notified until the problem is resolved. Alerts can keep problems from falling between the cracks.
2. A "buck passing" alert, which is triggered if a problem is passed around too many times.
3. A deadline alert. If the problem remains unresolved after a certain amount of time, an alert should be created no matter how hard someone has worked to solve it. Management can then decide whether to commit more resources.

Duplicate Checks

The ability to check for duplicate problem records by searching for duplicates of any field is very important. This can include searching on location, software, device type, serial number, etc. Opening two problems on the same item causes needless aggravation. It can also create confusion and dissipate resolution efforts if, for example, one problem is opened for a downed controller, and another is opened for a downed terminal connected to that controller.

Problem Management: Using the Help Desk

Reports

One of the major benefits of a well-designed problem management system is report generation. Management reports summarizing problem activity, mean-time-between-failure (MTBF), and mean-time-to-repair (MTTR) provide important decision-making support. Operational reports are important in providing continuity between shifts. More sophisticated tracking systems can provide inventory, maintenance, and configuration management information.

Availability tracking is also a major benefit. Network availability can be measured in many ways, as follows:

- *Explicit outage*, the percentage of time a component is out.
- *Implicit outage*, the amount of time a component is out because something it depends on is out.
- *Perceived outage*, outage which is only measured during operating hours.

The preceding information should be automatically and continuously updated as problems are encountered.

CONCLUSION

Establishing a Help Desk with a well-designed problem tracking system offers several benefits. It allows information to be integrated across a complex network and provides a historical view of network problems. This reduces the time between problem occurrence and problem resolution, increases system availability, and pinpoints emerging problem areas. It can boost users' confidence in the network.

Information and the ability to use it effectively are becoming a corporation's biggest asset. The network in which that information resides and across which it travels is key to the information's usefulness. A Help Desk is one essential part of managing that network efficiently. □

Designing a Practical Network Maintenance Strategy

This report will help you to:

- Reduce network downtime by developing a comprehensive maintenance plan.
 - Discover the hidden costs of neglecting user education in your organization.
 - Develop a new approach to maintenance which will save your company time and money.
-
-

Nowadays there is much talk of *managing technology*. But what is technology, and how does one manage it? How can a technology manager best supervise two fundamentally different elements—the technology and the human users? Certainly, attempts to manage what one fails to understand is an exercise in futility, particularly with respect to network management.

You can easily find out how little others, even your technical colleagues, understand about a network manager's job functions. Ask a wide sampling of technical, management, and clerical staff to define the term network, or describe it as they understand it. The most common explanation will resemble "*all of the wires and plugs that connect computers, telephones, etc.*" You will be disappointed with the quality of explanation, even from the technically adept staff—programmers, systems analysts, and even those who work on the network itself.

Because they fail to understand it, very few people will try to explain the purpose of the network and, most importantly, the human interaction that is central to the efficient working of a network.

This report was developed exclusively for Datapro by Anand V. Rao. Mr. Rao is president of Rao Communications, a technology reporting service based in New York City. Rao Communications specializes in technical documentation, market analysis, and reporting.

PUTTING THE NETWORK BACK IN CONTEXT

Networks do not exist in a vacuum. A network, however automated, connects human users. A network is a classic example of a technological facility interacting with humans. Even the most automated networks demand constant human supervision to make them work. Networks, like human beings, constantly fail to perform—for any number of reasons, and a successful network manager must be prepared to address errors from both quarters.

Managers of technology must contend with the fast shrinking distance between technical facilities and end users. Human buffers between them, such as telephone switchboard operators, have all but vanished. Nearly every office desk boasts a PC or terminal, and a modem-equipped telephone. These desktop systems grow ever more powerful, sophisticated, and easier to use. It no longer takes technical knowledge to use desktop electronic facilities. Most users have no idea how these devices work, nor do they much care—they simply appreciate their ease and convenience as tools for communication and *expect* them always to work. More and more users are connected to the networks, with access to centralized data bases and other computers, but few are trained in using the network communication facilities.

Designing a Practical Network Maintenance Strategy

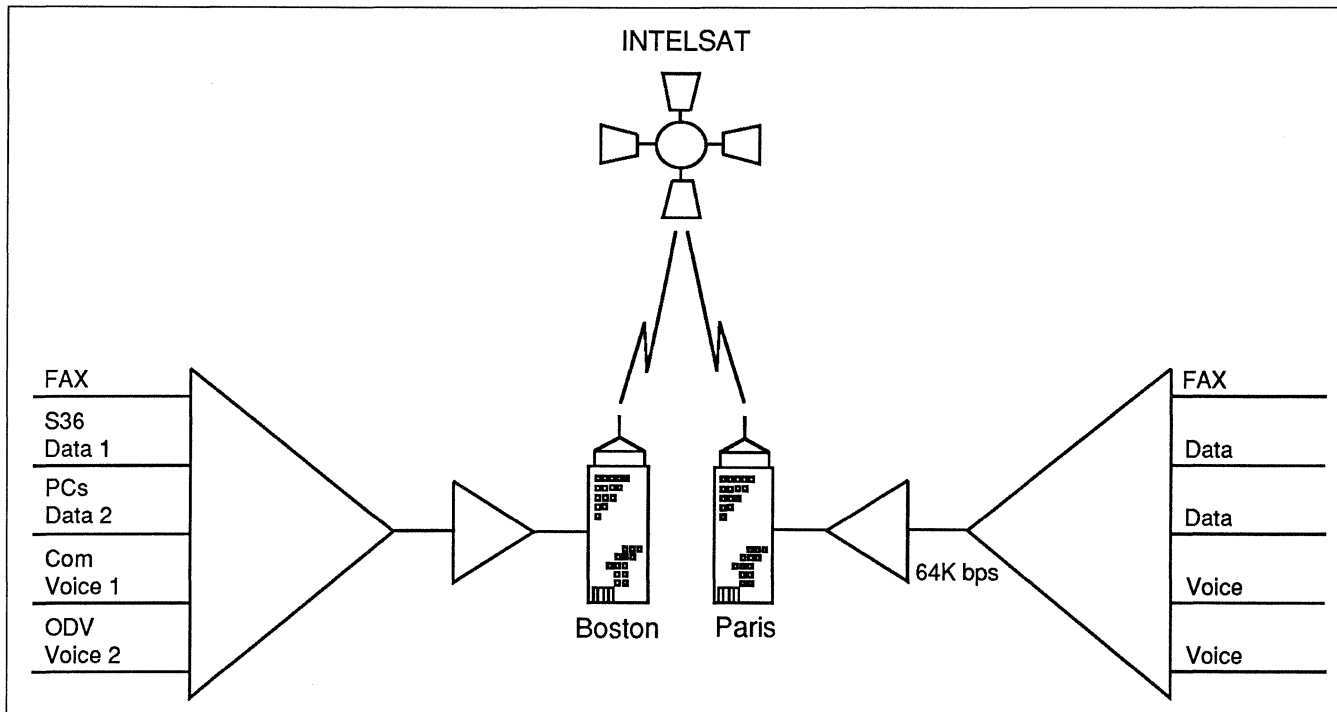


Figure 1. A simple diagram depicting a major link in the network of a large financial corporation. By developing documentation and diagrams such as this, the organization conveyed to its end users an overall picture of how the network functioned.

When trouble hits, users complain about the whole system and deny any fault. The network manager must anticipate that the system will fail, commit errors, and otherwise perform badly. Indeed, one of the primary job functions of the network manager is to prepare for network errors and failures. A successful maintenance plan involves providing suitable means to remedy errors, and ensures that the network continues to operate without interruption. Even the most sophisticated network management and control facility cannot overcome problems arising from a poorly planned maintenance strategy.

UP THE DOWN NETWORK: A CASE STUDY

This section describes the international network of a major financial organization. The experiences of the users in the organization are typical of those who interact with large networks.

The organization network consisted of several mainframe computers, front-end processors, multiplexers, voice/data integrators, electronic and voice-mail systems, links to outside facilities (provided by more than one phone company), and satellite linkups to offices in several countries. (See Figure 1.) Their network had so many problems that it seldom operated

as planned. Before logging on, most users invariably asked, "is the network up or down?"

When problems occurred, the network management staff huddled in conferences and hunched over control terminals trying to get the network working again. While the able network manager and maintenance staff usually restored service quickly, it was evident that the networks did not meet users' expectations.

The Users

The network users could be classified broadly into the following categories:

- Management staff who enjoyed some priorities and higher access levels.
- Clerical staff with lower access levels and lower priorities.
- Technical personnel, such as programmers, operations staff, and others.

The network was used to access several offices outside and inside the country by various levels of voice and data users. Also, it downloaded files used in

Designing a Practical Network Maintenance Strategy

running several programs on the company's main computers, located in its headquarters in another country.

An in-house survey of the network revealed:

- **Most users had no idea how the network functioned.** They had no idea what was connected to what, and how the desktop keyboarding that they performed enabled the network to perform the functions.
- **The operations staff did not know what happened to the data once it was transmitted out of the front-end processors.** For instance, very few of them knew the existence of a communications control room, or had any idea about the functions of the network maintenance personnel.
- **Upper management considered the functioning of the network too technical.** While upper managers approved \$50 million to build the network, their involvement in managing it was limited to complaining when it did not work.
- **Few users, even at highest levels, knew the overall functional objectives of the network.**
- No network diagram/graphics was available to inform the users how the entire network functioned.

Furthermore, the survey revealed no systematic error logging procedures, nor any operations documents. No training or error recover/trouble-shooting documentation was available. Moreover, no formalized procedures existed to report errors and network failures. Many users thought that the operations staff in the computer room was responsible for fixing problems. Others called programmers to complain about network problems.

Perhaps most important, the survey discovered that several network failures would have been prevented had users understood how the system functioned and been trained to use the network efficiently. While not necessarily typical of all networks, these survey results provide useful insights which can help prevent failures in all large corporate networks.

MAINTENANCE IS A STATE OF MIND

Most network users approach network maintenance as a *them vs us* issue. "Them" refers to technical staff who have the know-how and the responsibility to fix the problem. Many users don't even know how to describe the problem(s) they are experiencing, and yet expect the maintenance staff to know what they are

talking about. Others view maintenance as "it's your responsibility and you fix it."

Even the managers who authorize huge expenditures for installing the networks—influenced by their own staff, consultants, and vendors—think of maintenance as a "service contract" issue. Most managers view maintenance as problems that almost inevitably require vendor assistance. Managers using this approach feel that they have an adequate maintenance strategy as long as the vendor responds "yes" to questions such as, "Do you have a 24-hour service phone access?" and "Are your maintenance engineers available when we need them?"

Other users, including network managers, think of maintenance as a "fixing the fault" issue. They emphasize getting the system working again as soon as possible. Others view maintenance as a purely contractual question, and spend a lot of time arguing about dollars and cents issues in a typically incomprehensible contract document.

Very few believe that network management involves *everyone's* active participation, regardless of the user's access level and technical sophistication. Successful network maintenance requires that it be made an overall management concern—a policy matter and even a state of mind. It transcends the borders of narrow departmental concerns and the temptations to fix blame and pass the buck. Maintenance starts *before* the network fails and *before* it is installed; maintenance begins with the network plan.

THE HUMAN FACTOR: STILL THE VITAL ISSUE

Despite the capabilities of today's technology, maintenance and error detection involves much more than electronics to fix problems quickly. Human ingenuity and problem awareness remains the crucial factor.

El Al Airlines: Human Ingenuity at Work

For instance, most airlines and airports have installed sophisticated detection equipment to deter terrorists. Detection equipment is by no means foolproof, however. El Al Airlines, a prime terrorist target, does most of its baggage and passenger checking by trained security personnel. More than once, El Al security staff has discovered explosive materials hidden in luggage—even after the baggage passed electronic surveillance checks. El Al has the best safety record of all airlines and has never been hijacked or harmed by terrorists.

Designing a Practical Network Maintenance Strategy

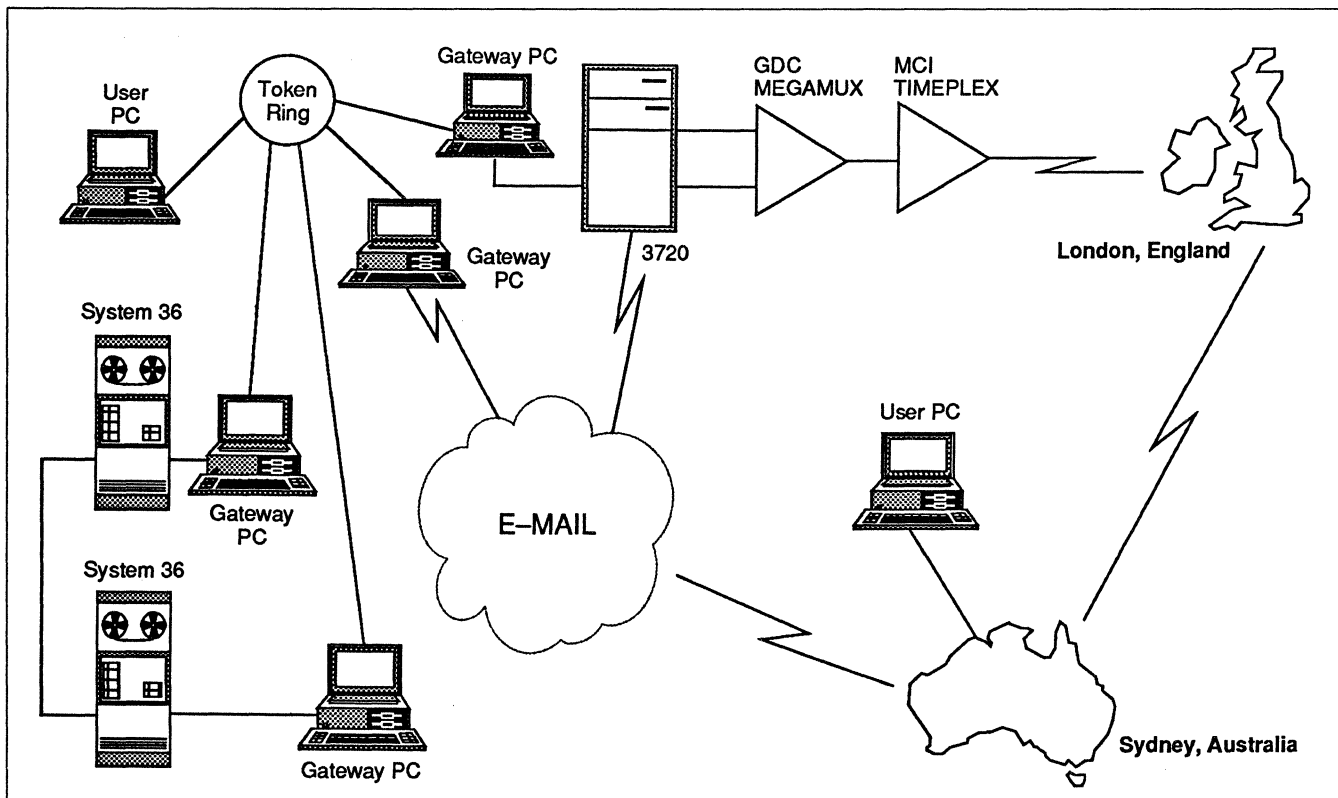


Figure 2. An electronic mail facility linking New York and Sydney, Australia. A simple diagram such as this can explain to operations staff the overall purpose and objectives of the network.

In formulating maintenance policies, network managers can benefit by emulating El Al's example. While network management and error detection tools need not be discarded, maintenance should emphasize training of maintenance staff and all network users. For example, El Al's security efforts are not limited to the personnel who check baggage and interview passengers before they board the flight. Every crew member is trained to spot potential trouble and report it as clearly as possible to those who can handle it.

While sophisticated test and error detection equipment is vital to successful network management, it should not displace the importance of human intervention and maintenance awareness.

Using the Tools

Operators need skills to decipher and analyze the displays of even the most sophisticated troubleshooting instruments. Quite often, the error messages are vague and difficult to understand. In many cases, the instrument indicates error conditions, but remedial measures require human analysis and judgement.

No two networks are the same; and qualified, trained maintenance personnel are always in short supply. Thus, it is very important to recruit suitable staff and train them on the network's specific requirements. Staff education should go beyond error detection and correction to include an explanation of the overall network's purpose and objectives. (See Figure 2.)

OPERATING STAFF: CRYING FOR RESPECT

Most operators in data processing and communications environments bemoan their lack of recognition in the organization. Although operators perform important tasks, organizations typically view operators as those who load, unload, and log. Often, the operators' talents and know-how are poorly used. Because the operator's role is often viewed as a dead-end job, the occupation suffers from high turnover and attrition rates.

Operators functions are often limited to performing routine work. Most are not required, and are often discouraged from, knowing anything beyond the immediate necessities of what they do. Many don't even

Designing a Practical Network Maintenance Strategy

know why they perform their routine chores and what happens in the network in the aftermath of their work.

Many operators are talented and can be trained to do some network maintenance tasks. With suitable training and recognition of their job's importance, operators can detect network failures. Operators should understand overall network functions, how the whole system connects, the importance of each unit, and their own role relative to the complete system. They should be trained and rewarded in network failure detection tasks, which can be performed without interfering with their routine work.

USER EDUCATION IN NETWORK MAINTENANCE

Although the network exists, for users, most have little idea how the system works. Perhaps users do not require an in-depth understanding of the system, but they should be trained not to commit the common mistakes that cause so many problems. Also, training can increase their awareness of the network. They can be trained in some dos and don'ts in using the system and learn the proper protocol of error reporting—how to describe the problem, and whom to report it to.

Even technically adept users, such as programmers, technical managers, and others, need help in knowing how to use the network. Amazingly, most organizations have no network user-manuals or materials that describe the networks' business purpose. Most companies don't even have a system diagram that shows how the various parts of the network are connected to one another.

THE NETWORK: A REVENUE CENTER?

In today's environment, where technology now sits on the office desk, managers generally view data processing and network facilities as direct cost centers. This is particularly surprising, considering that the very survival of many organizations and their ability to compete in today's marketplace depends on the quality of their data processing and network facilities.

Senior-level management expects an expensive network to perform well without much maintenance. Managers tend to purchase vendor maintenance contracts and expect that to solve all problems. Typically, senior management's attitude is "why spend more money in training users and operating personnel in maintenance procedures when the vendors are paid

to do it?" Network managers must persevere in convincing management that the money invested in training is well spent.

Making a Case for Training

The following points explain why training is important regardless of the money already invested in maintenance contracts:

- **Vendors supply only parts of the system.** They are not responsible for the entire network.
- **The network is configured according to the user's needs, not the vendor's.** As such, the problems are specific to the network.
- **Most network failures and consequent damages are due to incorrect usage and lack of trained personnel.**
- **It costs more than you think when the network is down for an hour.** Make a rough estimate of the cost of network downtime. While this will be an educated guess, at best, it will provide good material to shock management.
- **Your company can save money by implementing a comprehensive maintenance strategy.** Explain to upper management how instances of downtime could have been avoided, had users known how to detect errors and potential problems. Translate the downtime into dollars and project out how much the company can save by instituting a user-training program as part of a comprehensive network maintenance strategy.

Preventive maintenance costs little compared with the cost of fixing the network after it fails.

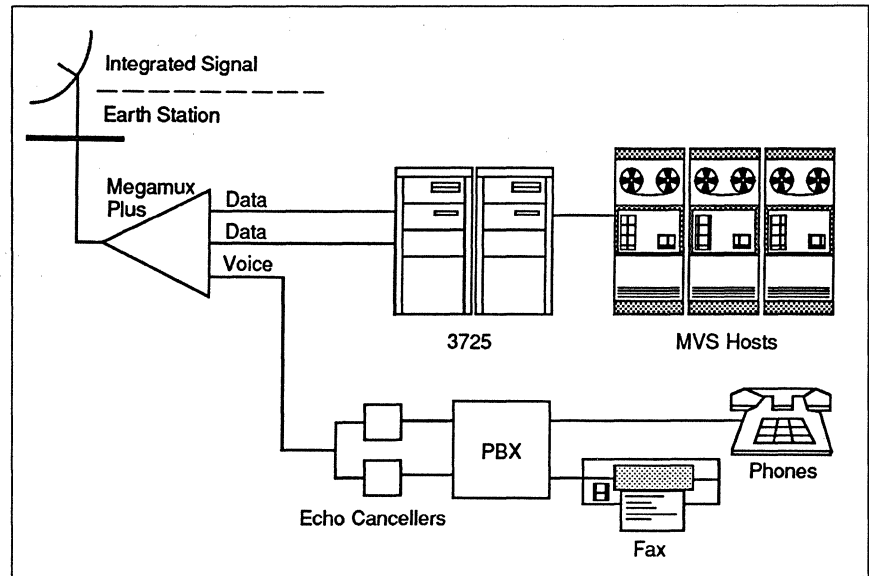
ISOLATING THE PROBLEM—NOT THE USER

The network manager must locate a problem before attempting to solve it. Locating the problem involves isolating it. Successful network maintenance strategies must provide the means to locate and isolate problems, as well as the facilities to fix them as soon as possible.

Users trained to understand the importance of their role in network maintenance can help network managers tremendously in locating and isolating problems. On the other hand, isolating users from understanding and participating in network objectives is counter-productive. Users must be trained

Designing a Practical Network Maintenance Strategy

Figure 3. Diagrams such as this can show the user how the network is connected. Illustrations need not be elaborate or technical; just sufficient to acquaint the user with the system, like a good map.



and made aware of their responsibilities in the smooth functioning of the expensive facilities they use.

User education and involvement goes a long way in maintaining the network. A comprehensive user training course should:

- **Educate the user on the overall functioning of the network.** This includes what is connected to what, how the network achieves its objectives, and how the system translates the user's keyboard input into objectives. This information need not be very elaborate or technical; just sufficient to acquaint users with the system like a good map. (See Figure 3.)
- **Train the user in using the desktop facilities to access the network.** Quite often, the users don't really know why they should enter all the required information during logon and other procedures.
- **Provide the user with manuals and other tools to use the system.** Update the manual as soon as changes occur in the network.
- **Train the user to clearly describe the problems they experience.** Provide users with a number that they can call for help.
- **Establish error reporting procedures.** Include steps for error logging and problem resolution.
- **Conduct periodic "system acquaintance and training" sessions to inform the users of the changes in the network.**

MONITORING NETWORK PERFORMANCE: A COMPONENT OF THE MAINTENANCE STRATEGY

It is amazing how much networks are taken for granted—most maintenance is patchwork. Little effort is made to tune up and improve network performance. Few organizations maintain a network performance log or any other formalized record.

Numerous network management tools on the market can display and log network performance. To be effective, however, they must support the overall maintenance strategy. Quite often, such tools are merely used when the network shows signs of trouble—rather than continually monitoring overall performance.

Often, staff members specialize in certain maintenance functions and may even discourage others from participating in their work. Even in large networks, it is not unusual to find only one or two technicians who can operate a certain network monitoring device. Thus, if the network fails, the absence of one specialist on a given day may mean that the fault cannot be fixed, especially if the organization maintains little or no documentation on the total working of the network. Monitoring a network without adequate documentation, graphics, and materials that describe the functional and business objectives of the system, is counterproductive.

Regular network monitoring must be made a part of the maintenance procedure and should include:

Designing a Practical Network Maintenance Strategy

- **Monitoring the physical status of network components.** This includes the hardware, transmission facilities, interfaces, cables, etc. It is absolutely necessary to keep a log on their conditions.
- **Monitoring performance.** This includes network use, traffic volume, response time, line and interface availability, capacity, etc. Once again, regular logs should be maintained.
- **Keeping track of user complaints.** User complaints reveal shortcomings in user education and, quite often, the inadequacies of the network itself.
- **Monitoring the patterns of network problems.** Networks often are prone to persistent problems within certain sections of the system.
- **Monitoring vendor maintenance performance.** It is essential to monitor and log all vendor maintenance performance.
- **Monitoring personnel performance.** This can be done without putting them under unnecessary pressure.
- **Periodic user and network manager meetings.** A lot can be achieved in these encounters.

The financial organization described previously in the "CASE STUDY" instituted all the above monitoring functions, and thereby eliminated several of its constant problems.

HELP DESK OR COMPLAINTS KIOSK?

The help desk staff is often placed there on a short term basis, and members view their situation as a temporary assignment. The staff tends to hide behind jargon and offers users explanations that are difficult to understand. Help desks should be operated by trained personnel who have the patience and the right attitude for the job.

In many establishments, help desks do not log calls. The importance of logging can not be overstated. The quality, frequency, type, and repetitiveness of help desk calls reveal a lot about network problems.

Network managers can use help desk logs to plan for training, maintenance, and other management functions. Additionally, the help desk is an essential part of the network's overall maintenance strategy. A centralized help desk is vital to the efficient functioning of the network.

Network maintenance is often beset with problems merely because the work is considered a chore. Often, maintenance is performed as needed, rather than as a long-term, planned function. A well-conceived, documented, and judiciously implemented maintenance plan can contribute significantly to the overall management strategy of the network. Ignoring or relegating maintenance to secondary status can prove costly. □

Future Scenarios for Network Management

This report will help you to:

- Understand the concept of integrated network management.
 - Examine the relative importance of IBM's NetView, AT&T's UNMA, and Digital's EMA.
 - Anticipate future trends and scenarios for network management systems.
-
-

Integrated network management is a new definition of network architecture. Although there is currently more smoke than substance regarding the topic, IBM has raised upper management's awareness to its criticality and to the increasingly serious network management situation which exists today. IBM flew in upper managers, many of whom were from the largest corporations, and convinced them that network management was a major strategic issue. Correct implementation could lead to increased productivity, availability of information, control of service dollars and procurements through ongoing capacity planning, and many other familiar benefits.

IBM has also raised upper management's expectations. However, upon management's urging, MIS directors inquired about early delivery of NetView—only to be shown a view graph by IBM representatives. That situation has presumably been at least partially remedied with recent product announcements.

Network management is not a new concept. Network management systems have been commonly called accounting systems, network diagnostic tools, and various other names. The current experience is that networks change quite a bit over time. Network archi-

tectures are no longer homogeneous; i.e., almost every network in the industry is a hybrid of local area networks (LANs) tying into Systems Network Architecture (SNA) or X.25 backbones. There are visions of integrating voice and data as well as transmission management. Existing network management products do not readily tie these different systems together. As a result, it has been virtually impossible to get some kind of end-to-end picture of a modern network. This picture may be changing, however.

Is NetView or any other network management system from AT&T, DEC, Hewlett-Packard, or an RBOC going to provide the tools to perform the network management functions listed in Table 1? What is an integrated network management architecture and do corporations really need it?

FOCUS ON PEOPLE

The tendency has been to focus on the technology aspects of integrated network management. The real network management issue is the lack of adequately trained people. The hope is that technology will solve what is fundamentally a staffing issue. In network control centers throughout the country there are operators sitting at separate control consoles for different disciplines such as transmission, switching, data circuit-terminating equipment (DCE), SNA, LANs, and other types of networks. Throughout the world, networks are moving more toward 24-hour-per-day,

This Datapro report is based on "Future Scenarios for Network Management," by Vince Barrett, Ernst & Whinney, from *Telecommunications*, January 1989. © 1989, Horizon House-Microwave Inc. Reprinted by permission.

Future Scenarios for Network Management

<u>The Traditional Areas</u>	
<p>Problem Management</p> <ul style="list-style-type: none"> • Problem determination • Problem diagnosis • Problem bypass and recovery • Problem resolution • Problem tracking and control <p>Performance Accounting Management</p> <ul style="list-style-type: none"> • Usage • Responsiveness • Availability • Cost 	<p>Change Management</p> <ul style="list-style-type: none"> • Additions • Deletions • Modifications <p>Configuration Management</p> <ul style="list-style-type: none"> • Logical resources • Physical resources • Relationships
<u>The New Areas</u>	
<p>Systems Management</p> <ul style="list-style-type: none"> • "Box" level problem management • Remote operations • Software management • "Lights Out" management 	<p>Asset Management</p> <ul style="list-style-type: none"> • Capacity planning • Network modeling • Overflow management • "Least Cost" management
<u>The "Emerging" Areas</u>	
<p>Security Management</p> <ul style="list-style-type: none"> • Disaster recovery/contingency management 	<p>Directory Management</p> <ul style="list-style-type: none"> • LAN management

Table 1. Network management subsets.

7-day-per-week, and 365-day-per-year operations. In the securities industry, for example, the idea of linking all of the exchanges in the world is making the information arriving at night just as critical as the information received during normal business hours.

The problem during the graveyard shift is that the transmission facility in use can be less reliable than the one used during the main shift. This means that the best people must be in the network control room during off-hours. The issue now becomes one of cost to manage the network. Cost containment is driving the industry to explore technologies such as artificial intelligence (AI) and integrated network management systems. The need is to develop a meaningful common verbal language for use by a reduced population of console operators.

Demand for people will depend on architecture. There is no real need to have a single integrated network management data base because a single integrated network management architecture will not exist in most companies. The requirements of most companies today are much broader than the capabilities available from vendor-specific network management systems on the market now or in the near future. Telecommunication managers should strive for interoperable data bases when putting together network management systems. Although there are no

strategic products in this industry today, an interface like Open Systems Interconnection (OSI) becomes very important.

While many organizations may want a "seamless" network management system, most will opt for functional integration. The absence of a network guru or a full spectrum of networking technicians will force organizations to adopt a subsystem approach to network management. For at least the next 3 to 5 years, until meaningful relief comes from artificial intelligence, network control center staffing will continue to be a more critical problem than technological shortfalls. Although AI-based solutions offer long-term promise, they will be slow in developing and low in functionality for the next 36 months. The two most critical pieces of the network management architecture will be syntax consistency and interoperability between network management data bases.

CENTERS OF GRAVITY

There are two centers of gravity in the universe of network management systems. One center of gravity can be represented by NetView while the other is represented by OSI and a group of vendors like AT&T (with UNMA), Tandem Computer, Hewlett-Packard in conjunction with Northern Telecom, and DEC. In between these two centers are all of the communication vendors being pulled in both direc-

Future Scenarios for Network Management

tions. In such a bipolar universe, a vendor either supports NetView or OSI in the same way that it supports either SNA or OSI.

An embarrassment of the OSI community may be the inability of OSI-based network management systems to interoperate. Some standards are needed here. The problem with standards is that an infinite number of ways exist to implement them—none of which ensure compatibility. An encouraging sign, however, is the recent announcement of the formation of an OSI network management committee involving many of these major vendors, i.e., AT&T, Hewlett-Packard, and Northern Telecom. Ironically, it is this area of network management chaos which is the most likely to show the benefits of the standards. Fear of NetView will probably bring about a standard graphics-based “man-machine” network management interface.

Another visible phenomenon is the increasing use of “brochure compatibility” in the industry. Because vendors have a limited amount of development resources, many are marketing their ability to be compatible with UNMA, NetView, and OpenView. It is not possible for manufacturers to put all of these capabilities into their products until clear winners are determined.

PEER-TO-PEER NETWORK MANAGEMENT

Will NetView be open at the Focal Point Level? NetView is an open network management architecture that is open in the same way that SNA is an open network architecture—that is, as long as there is 3270 emulation. The NetView/PC is really similar to a 3270 device—it is still a master/slave management system. Minicomputer-based systems, T1 multiplexers, PBXs, and the new T3 multiplexers entering the market will soon require a peer-to-peer level management for passing all of the network management data. Currently, they are prevented from doing so. The myth of openness is not just restricted to NetView. It is also a myth in the OSI-based network management architectures. Open networks is a concept that is really in the eye of the beholder.

Non-IBM system vendors will ultimately force IBM to open NetView, but not until 1990, stimulated by possible diversion of user-buying dollars. Non-IBM system vendors will be forced to work toward common OSI-based network management standards, however, and the adoption of this standard will be painfully slow. IBM will be able to capture the critical area of network management data base, but in opening NetView, will find it difficult to lock down network management. In other words, NetView will not

be sufficient and timely enough to serve as the integrated network management architecture.

Optimistically, telecommunication managers will be able to run their networks with three separate integrated network management systems. IBM customers probably will have NetView, and DEC customers will most likely have EMA. In addition, there will be the requirement to manage the PBXs through a carrier-based network management system. The best scenario for the near future is that three separate network management architectures will need to be tied together for interoperability (not integration) of the network management data bases and through common user interfaces.

REAL PROBLEMS, REAL ANSWERS

What do telecommunication managers do now? The problem is that the strategic planning decisions will require that managers make tactical decisions with strategic implications. Managers must be able to look ahead and determine the direction of the technology being used and ask themselves, “Do I have to throw away this system in 24 months if a vendor like IBM indicates that it will announce a similar system in the near future?”

There are fundamental decisions that managers must make to deal with the complexities of highly stratified networks. As networks in corporations become more like profit centers, network managers will be administering accounting packages, performance monitoring equipment, network modeling tools, test equipment, matrix switches, consoles, and data bases—the fundamentals of a network control center. These may not be considered strategic, but all of this equipment is necessary for making a network operate efficiently.

The strategic direction versus tactical necessity will force users to make third-party network management buying decisions. The network management vendor marketplace will prove to be more volatile than the LAN market—leaving a network management “trail of tears.” Network planners should fill higher-level network management functions reluctantly for the next 24 to 36 months. Test and delay and delay as long as possible. This is a time when “analysis paralysis is your friend.” Finally, telecommunication managers can focus on the fundamentals by building a strong, sound network control center infrastructure.

In the absence of viable integrated management products, AI-based product breakthroughs, and full spectrum networking, super technicians will force organizations to adopt a subsystem approach to network management. Network control center staffing

Future Scenarios for Network Management

will prove to be a more critical problem than technological shortfalls in the next 2 or 3 years.

NetView will be adopted by many organizations as the strategic network management direction; however, tactical necessity will severely dilute the implementation of that direction. DCE and LAN management will be hotly contested areas for network management between NetView and OSI-based solutions. The system vendors, but mostly users, will ultimately force IBM to open NetView at the Focal Point Level.

The existence of NetView/PC does not mean that the network is open because the next level also has to be open; peer-to-peer level network management is an absolute requirement for integrated network management. NetView will fail to capture PBX/voice network management. In the next few years, there will be no easy decisions. Once again, the users will play the role of systems integrator. □

Using Network Management Systems to Gain a Strategic Competitive Advantage

This report will help you to:

- Learn how major corporations use network management systems to increase and sustain their competitive advantage.
 - Use the capabilities of your existing network management equipment to maximize your network's strategic value.
-
-

Traditionally, users viewed corporate networks as mere utilities that supported the organizational infrastructure by facilitating communications among far-flung locations. The control inherent in private facilities translated into substantial cost savings over the long term—thus justifying the expense of running one's own network.

In the early 1980s, industry experts promoted private networks as the means of making geographically dispersed companies more manageable—achieving cohesiveness among diverse operating units and helping top-heavy organizations trim the burgeoning ranks of middle management. Ideally, a private network would promote better and more timely decision making, translating into a more profitable business. The ideas had some merit, but for a long time the appeal of private networks remained focused on anticipated cost savings and on increasing the flexibility in allocating communications resources.

Today, a ground swell of opinion says that the quality of a company's network holds the key to serving customers better, increasing market share, and pursuing new business opportunities successfully. In the process, an enterprise can secure strategic competitive advantages over marketplace rivals that have not yet awakened to such possibilities.

This report was developed exclusively for Datapro by Nathan J. Muller. A former consultant, Mr. Muller has 18 years experience in the computer and telecommunications industries. He has written extensively on all aspects of computers and communications and is the author of "Minimum Risk Strategy for Acquiring Communications Equipment and Services" (Artech House, 1989).


Many accept this vision of the corporate network as a competitive weapon as self-evident. Without further explanation, however, one might be led to believe that the particular arrangement of lines or the type of equipment deployed among the various nodes determines the network's strategic value.

Although certainly essential, network architecture and components take a backseat to the network management system (NMS). The right NMS unifies diverse computer and communications resources and transforms them into strategic assets that improve a company's competitive position and long-term survivability. Not surprisingly, then, many organizations place a premium value on network downtime:

- A Wall Street brokerage house can lose as much as \$60,000 per minute when buy/sell instructions from customers are disrupted.
- A state lottery can lose millions of dollars per hour if it cannot process ticket sales when hefty jackpots are at stake.
- An insurance company can lose its *Fortune* 500 accounts if it cannot live up to specified levels of network uptime to process its clients' claims.

As more and more companies are discovering, the capability to maintain a strategic competitive advantage rests upon the quality of their networks. The "network" includes the high-capacity backbone, feeder links, and drops typically associated with wide area networks (WANs) as well as the host, servers,

Using Network Management Systems to Gain a Strategic Competitive Advantage



The Texas Air network will use cross-connect systems and multiplexers from DSC Communications Corp. This photo shows a DSC DEX CP2000 Digital Network Access System, a byte-interleaved multiplexer that provides cross-connecting at the DS0 or DS1 level. Monitoring and control are enhanced by connection with the DSC DEX NMS management system (see next photo).

and terminals of local area networks (LANs). Not only the links and the hardware must be monitored for proper operation to ensure maximum network availability, but the information traversing the links and the integrity of the applications must also be protected. Only a comprehensive management system can ensure network quality by providing a view of the entire network, extending diagnostics and control to the farthest corners.

CASE STUDIES

Transportation

Among the largest private voice-data networks under construction is that of Houston-based Texas Air, the holding company for Eastern and Continental Air-

lines. The network—System One—will include 20,000 miles of digital lines, tying nearly 50 cities into a 10,000-mile, high-speed fiber optic backbone provided by LightNet (Chevy Chase, Maryland).

The main motivation for a heavy investment in a private network is the high cost of the public telephone network. Prior to setting up its own network, Texas Air was AT&T's seventh largest customer; its two airline companies were spending \$200 million annually just for handling flight reservations.

When completed in 1995, approximately 115,000 terminals and microcomputers will connect to Texas Air's System One network, providing access to reservation centers and travel agencies nationwide. The network will use cross-connect systems and multiplexers from DSC Communications Corp. (Santa Clara, California) to collect and transport voice and data communications, as well as to provide access to IBM, Amdahl, and Unisys mainframes at network control points in major cities. Separate management systems control the hosts, whereas DSC's DEX NMS series of management systems will control the wide area network. In addition to handling remote monitoring, the DEX NMS management systems may be used for automatically rerouting traffic around failed lines and reconfiguring the network to bypass failed equipment. Texas Air projects a \$200 million savings by the time System One is fully implemented.

The network's strategic importance to Texas Air is quite substantial. Every call received over the network translates into a gain of about \$27, whereas each call missed translates into a \$6 loss. The restoral capabilities of the DEX NMS systems will ensure that the causes of lost calls are virtually eliminated. With an anticipated traffic load of 15 million voice and 800 million data calls each month, System One might very well spell the difference between profit and loss in an industry that operates with very tight margins.

Banking

Like airlines, banks must reach out to customers to stay competitive. The long-term objective of Pennbancorp (Titusville, Pennsylvania) in an industry that is steadily shrinking through consolidation, is not merely to survive, but to grow. Rather than adding branches, Pennbancorp's growth strategy emphasizes acquiring other banks as the most economical way to increase market share.

After the breakup of AT&T and the subsequent fragmentation of the telecommunications industry, Pennbancorp realized that it could not fully integrate the

Using Network Management Systems to Gain a Strategic Competitive Advantage

operations of other banks without first unifying and controlling its own network.

Pennbancorp began its push to unity in 1985 by merging the five networks of its three banks into a single high-speed backbone network. The backbone consolidates traffic from various regional branches, yielding substantial savings in long-haul transmission costs and, in turn, accelerating the payback on the hardware required to implement the network.

To implement its high-speed digital backbone network, the bank deployed five medium- and three high-speed multiplexers purchased from General DataComm (Middlebury, Connecticut). The five Megamux Plus multiplexers act as "feeders," consolidating traffic from various locations for economical long-haul transport to the three hub locations served by GDC's Megaswitch nodal multiplexers. The network is centrally managed from the Megaswitch controller located at Pennbancorp's Oil City network control center.


To support its trust network, Pennbancorp put into operation 15 GDC statistical time-division multiplexers, which provide centralized diagnostic and control as well as drop, insert, and bypass capabilities. Pennbancorp also deployed 130 GDC diagnostic modems which have enhanced the bank's capability to manage the network down to every drop location. The modem network is tied to GDC's NETCON-70 Network Management Controller, which provides centralized diagnostic and control features. Dial backup procedures are implemented automatically to reroute data over the public switched network.

The combination of equipment and lines—plus the efficiencies gained from the NMS of each type of device—has already shaved about \$500,000 annually from Pennbancorp's operating costs, while giving it the flexibility to deal with a wide variety of equipment and services acquired through the acquisition of other banks.

The Pennbancorp network now consists of 200 branch offices and 87 automated teller machines (ATMs) sprawled across the western portion of the state, from Erie in the north to Uniontown in the south. Its capability to accommodate the networks of newly acquired banks helped lift Pennbancorp into third place among banks in western Pennsylvania. It now controls \$3 billion in assets.

Air Package Delivery

Networks can also act as the vehicle through which companies enter new markets. United Parcel Service



The DSC DEX NMS-L management system uses a Sun Workstation (shown above) to provide a consolidated graphics-based interface to its DEX CP2000/1000 multiplexers and DEX CD digital cross-connect products. Texas Air will use DEX NMS management systems to control its wide area network, scheduled for completion in 1995.

(UPS) recently expanded its international air package delivery service into 19 more countries, bringing the total number served to 41. Such an intrepid move would not have been possible without the company's two-year-old international network.

The UPS International Shipments Processing System (ISPS) greatly reduces the delay in international deliveries by generating the documents needed to clear packages through customs. The documents can be transmitted directly to the computers of customs-houses or customs brokerage houses serving ports of entry. This allows UPS to begin the clearance process hours before packages arrive. In this way, the 550,000 packages shipped by UPS to international locations daily can be delivered to addressees with minimal delay.

At the company's main data center in Paramus, New Jersey, three IBM mainframes support package-handling operations around the world. Hundreds of microcomputers tie into the mainframes via IBM terminal controllers, which concentrate the traffic from various locations onto high-speed leased lines—including T1 and X.25 packet switched services. UPS thus can track international air packages from the point of exit in the exporting country through the customs operations of the importing country.

UPS is well along on a five-year plan that will transform its network into the transportation industry's most advanced communications operation by 1991. The \$1.4 billion network upgrade, which includes several million dollars earmarked for diverse network management systems, is expected to give the company a commanding lead over the competition in

Using Network Management Systems to Gain a Strategic Competitive Advantage

what UPS predicts will be a thriving business. The market for international delivery service is still in its infancy.

Companies of all types and sizes now view the way they organize their communications resources as a means to gain competitive advantages. Consequently, such strategic resources cannot be allowed to fail; nor can delays resulting from traffic overloads be tolerated. Given the increasing reliance of corporations on their networks, the task of keeping them up and running has assumed a top priority among technically savvy executives, who find network management systems indispensable tools.

THE WAY THINGS WERE

Only a decade ago, network management systems didn't exist. Building networks and getting them to work properly were the real priorities and, by today's standards, corporate requirements for voice and data communications were relatively simple (see Table 1). Telephone service was typically provided through a PBX or Centrex. Any "management" that had to be done was taken care of by the local telephone company. Aside from job scheduling and preventive maintenance on hardware, computer resources didn't require much management, since they were typically used only by the company's MIS/Data Processing group. Microcomputers were just being introduced to the office, so there wasn't a big demand for LANs.

The introduction of station message detail recording (SMDR) as an add-on to PBXs proved to be a milestone in the management of voice networks. With a service bureau or third-party software package, users could turn raw call data generated by the PBX into management reports detailing usage and costs, which could be used to support decision making and planning.

On the MIS side, users contended for mainframe access via a front-end switch. Modems were required to access remote computers. Ordinary telephone lines provided links for low-speed data transfer, but at the risk of errors from a variety of voice frequency impairments. Dedicated lines with conditioning provided higher quality transmissions at faster rates. The intelligent modems developed in the 1970s for use on leased lines offered rudimentary diagnostics and were used for basic network testing.

Later, centralized management and control systems appeared, allowing users to monitor the entire modem network from a single location. By the mid-1980s, vendors were introducing more sophisticated management capabilities and integrating them into

many other types of products, such as multiplexers, cross-connects, and LANs. These network management systems give technicians the ability to diagnose and correct problems anywhere on the network. Additionally, the systems can even predict the likelihood of problems, allowing operators to divert traffic from failing lines or equipment, with little or no inconvenience to users.

BASIC NETWORK MANAGEMENT FUNCTIONS

Although no universally accepted definition of "network management" exists, a consensus is developing about what such systems should include, driven, in part, by the increasing acceptance of OSI Management standards. (For more information, see "OSI-based Network Management," Report NM40-200-101.) The five main functions of network management are:

- Fault management (fault detection and isolation)
- Performance measurement
- Configuration management
- Security management
- Inventory/accounting management

Some industry experts also include two additional categories—maintenance tracking and applications management

Each of these functions contributes to enhancing the network's reliability and efficiency—and hence, its strategic value.

Fault Detection

With fault detection and isolation capabilities, users can find out whether problems are caused by equipment failures, line outages, or both.

Today's advanced NMSs detect problems by continuously monitoring performance and automatically conducting diagnostic tests on the variables that impair transmission. NMSs can be comprehensive in terms of diagnostic capabilities, performing a full repertoire of voice frequency (VF) impairment measurements as well as bit error rate tests (BERTs) on digital facilities. Some systems prompt the operator to answer questions and then recommend diagnostic tests to perform.

Using Network Management Systems to Gain a Strategic Competitive Advantage

NMSs generate alarms to report equipment malfunctions and prioritize them to expedite resolution. Some systems can detect such problems as modem power recovery, streaming terminals, corrupt configurations, and front-panel tampering.

Color displays make it easier for the NMS console operator to spot a problem and instantly obtain a reading on its severity through the use of multicolor alarm indicators. Various alarms can describe network problems and pinpoint the sources of those problems so that alternate facilities can be manually or automatically substituted.

Maintenance Tracking

Maintenance tracking relates to detection and isolation. Maintenance tracking is accomplished by using a database that accumulates trouble ticket information. A trouble ticket records the date and time a problem occurred, the specific devices and facilities involved, the accountable vendor, the name of the operator who responded to the alarm, and any short-term actions taken to resolve the problem.

Network managers can use this information for long-term planning and decision support. The manager can review reports that show failures of particular network segments, failures of a specific type of device or vendor's product, and problems that have not been resolved within a specific time frame.

Some NMSs automatically open a trouble ticket and send it to the network location that can act on it. If the problem is on a leased circuit, the trouble ticket is automatically sent to the appropriate carrier. Multiplexers from different manufacturers that have an interface to NetView/PC can support each other by routing trouble tickets through the interface.

Performance Measurement

Performance measurement has two aspects: response time and network availability. Many network management systems measure response time at the local end, from the time the monitoring unit receives the "start of transmission" command or "end of transmission" command from a given unit. Other systems can measure end-to-end response time at the remote unit. In either case, the network management system displays and records response time information and can generate user-specified response time statistics for a particular terminal, line, network segment, or the entire network.

Response time information may be reported in real time from the management terminal, which provides user-friendly icons and menus. Such systems display elaborate network schematics with specific colors assigned to various levels of response time. These displays can be used to identify the cause of a delay. When an application overruns its allotted response time, the manager can decide to reallocate terminals, place more restrictions on access, or install faster equipment to improve response time.

With on-line transaction services, such as those provided by automated teller machines, airline reservation systems, and point-of-sale credit card verification terminals, keeping response time within acceptable limits is essential to revenue generation. Too much delay causes impatience, and customers tend to differentiate service providers on the speed with which they deliver information. Without an NMS capable of measuring delay, spotting problems in a timely manner and ascertaining the cause of the delay is virtually impossible.

For example, delay is sometimes the fault of the carrier, which may be rerouting traffic over longer distances to balance its network load. With the ability to measure internodal delay through the NMS, operations personnel are in a position to request a better route from the carrier and to monitor the result of that change. If delay problems persist, internal causes can be investigated and traced to a congested node, degrading link, or faulty equipment.

Availability is a measure of actual network uptime, either as a whole or by network segments. Managers can compile statistical reports that summarize total hours available over time, average hours available within a specified time, and meantime-between-failure (MTBF). Some network management systems allow users to customize the report formats to eliminate reams of redundant information and to combine graphics with columnar data, making the information more comprehensible. Availability reports facilitate planning, enabling the enterprise to secure and maintain its strategic competitive advantage.

Long-term response time and availability statistics, compiled and formatted by the network management system, provide managers with objective tools for uncovering current trends in network usage, predicting future trends, and planning the assignment of resources for specific present and future applications.

An alternative to a network management system is to depend on computer vendors and carriers to find and correct problems. But during the interim, an organization's productivity suffers and revenue losses

Using Network Management Systems to Gain a Strategic Competitive Advantage

mount, if only because employees must be paid for waiting until service is restored. Additional losses may result from lost business opportunities, poor customer service, and ill will created among suppliers and creditors.

Configuration Management

Network management systems also provide the means to configure lines at remote locations. If a line becomes too noisy to handle data reliably, for example, the system will automatically reroute traffic to another line or the public network. When the quality of the failed line improves, the system will reinstate the original configuration.

Equipment, too, may be configured remotely. Features and transmission speeds of software-controlled modems may be changed. If a nodal multiplexer fails, the management system can call redundant components into action or invoke an alternate configuration. And when a node is added to the network, the management system can devise the best routing plan for the traffic it will handle.

Applications Management

Applications management is the capability to alter circuit routing and bandwidth availability to accommodate applications that change by time of day. Voice traffic, for example, tends to diminish after normal business hours, while data traffic may change from transaction-based to wide-band applications such as inventory updates and large printing tasks.

Applications management includes the capability to change the interface definition of a circuit so that one circuit can alternately support asynchronous and synchronous data applications. It also includes the capability to determine appropriate data rates to meet the response time objectives of various applications or to conserve bandwidth during periods of high demand.

Such network management capabilities are especially important for financial institutions that depend on their networks to place customer orders, transfer funds, and provide decision-support information. The stock transactions of brokerage houses on Wall Street, for example, are usually heaviest between the hours of 3 and 4 p.m. In such an environment, the capability of the NMS to downspeed circuits automatically by time of day means that two to four times as many calls can be handled without adding lines.

Security

Network management systems now address the security concerns of users. Terminals used for network management can be password protected to minimize disruptions due to database tampering. Various levels of user access can be set to prevent accidental damage. Senior technicians, for example, can have passwords that allow them to change the various databases, whereas less experienced technicians' passwords would allow them to review databases but not make changes.

Individual users, too, can be given passwords that permit them to make use of certain network resources, but deny them access to others. Methods also are available to protect networks from intruders who try to dial into computers with modems.

Inventory/Accounting

A network management system also allows users to keep an inventory of the network, including the number and types of lines serving various locations, and the capabilities that exist for alternative routing. In fact, the CRT screen of some NMS terminals can depict an international network map, using multiple colors to show equipment and line status. The operator can view individual network nodes and zoom all the way down to the actual card arrangements in the equipment cabinets. Some systems track the depreciation of components to facilitate corporate accounting. Managers can use this capability to identify underutilized lines, possibly eliminating the need to order more circuits, and to locate spare circuit boards, terminals, modems and other network components for possible reallocation, thereby reducing unnecessary purchases and boosting network cost efficiencies.

Using NMSs to gain strategic competitive advantage entails more than exercising close administrative control over the network. As mentioned before, the NMS must also provide the means to spot potential problems before they degrade the performance of the entire network. The NMS must be capable of initiating the rapid recovery of network segments through rerouting, providing the means to model the network to test various disaster recovery scenarios, and when necessary, handing off control from the failing master node to a subordinate node. □

Eight Critical Steps in Evaluating and Implementing Network Management Systems

This report will help you to:

- Determine your organization's specific network management needs.
 - Assess vendor network management systems in terms of your organization's needs.
 - Develop a general plan for implementing a new network management system.
-
-

STEP ONE—NEEDS ANALYSIS

The first step in evaluating a network management system is to document the current environment for all relevant application areas. This includes outlining procedures and job descriptions, as well as documenting network status reports and other information that specific individuals or systems depend upon. It also includes documenting work loads and problems associated with the current modus operandi. In addition, it is important in all areas to project out these needs for three years, in order to identify potential problems and future plans.

One method of gathering this information is to meet with key personnel responsible for the following areas:

- Voice Telecommunications—workorders, trouble desk, operators (for evaluating directory needs), network control center, technicians (for evaluating cable requirements).
- Data Communications—same as above.

This report is based on "Eight Critical Steps in Evaluating and Implementing Network Management Systems" by John Dretler, The Info Group. © 1988, The Info Group. Presented at the 1988 Network Management Solutions Conference, Boston, MA, June 18-21, 1988. Reprinted with permission.

- Finance—to determine cost allocation, departmental budgets, cost control, bill reconciliation, asset tracking.
- Human Resources—for evaluating directory information needs and staff responsibilities.
- Corporate Management—for setting quality of service goals and evaluating strategic opportunities.

Meeting participants should discuss their needs relative to each of the following functional application areas:

- Call Accounting—cost control, allocation, budget and input to network design.
- Inventory Control—cost allocation, asset management.
- Workorder Management—input to other modules, quality of service, vendor performance.
- Trouble Reporting—quality of service, system and vendor performance.
- Cable Management—cost control, availability.
- Directory Services—quality of service.
- Network Design—cost control.

Eight Critical Steps in Evaluating and Implementing Network Management Systems

- Trunk/PBX Monitoring—quality of service, cost control, network design.
- General Accounting/Bill Reconciliation—cost control, cost allocation, budget.

At the conclusion of these meetings, you should be able to determine the areas of concern shared by others in the organization—and from whom you can get support. Develop a written needs analysis to document your findings. (See the “Example Case Study”: at the end of this report.) After reviewing your findings, select the application areas that are most important to the organization. Next, analyze these areas in detail before commencing to Step Two—Product Assessment.

STEP TWO—PRODUCT ASSESSMENT—APPLICATION AREAS

The goal of Step Two is to assess how well a particular product (or products) performs relative to the critical application areas identified in STEP ONE. For example, one critical area could be 100 percent cost allocation. Another could be a directory—perhaps in twelve months, when an organizational change in Human Resources is planned. A third application area could be cabling—perhaps in three years, when a new PBX and cabling plan will be implemented. A fourth application area may be the eventual centralization of telecom management from separate locations. Again, it is critical to establish a plan that has a time line of approximately three years.

Use a three-year schedule to take the first and most critical step in product assessment: Define the data types which the product must support in order to meet the application and information objectives of the organization. (See Table 1 for a sample data matrix.) Evaluate the product's ability to:

- Support all required fields of information—for example, if you want equipment inventory, do you want all features identified? If you want cost allocation of workorder, do you want to have multiple types of charges? etc.
- Provide field sizes that are large enough for your information needs—for example, examine the product's field sizes for organization code (and the number of levels) and the size of the name field in directory. Also, determine whether or not the cable numbering scheme is compatible.
- Handle the varying types of transactions required—Does the product provide the proper cost allocation

on net/offnet calling? Does it handle multiple workorders for a bulk move? Can it vary directory lookup, etc.

If the product is compatible from an information standpoint, next examine its *usability*. Is it easy to learn initially? Can the user “express” through the system once learned? (This will prevent system use from becoming tedious). Is there a single point-of-entry for all information? Basically all data entry for all modules should be done via the workorder. When this is closed, all files should be updated.

Finally, the product should be modular, allowing the user to implement selected modules on a time schedule that suits the organization.

STEP THREE—PRODUCT ASSESSMENT: TECHNICAL CAPABILITIES

The user must consider a number of technical criteria when evaluating a network management system. As in Step One, it is critical to consider the organization's needs for a three-year horizon.

When evaluating a potential network management system against technical criteria, it is important, and also practical, to keep in mind the corporate computing philosophy—e.g., centralized, departmental, IBM, Digital, UNIX, PC, mixed, etc. Most managers find a network management solution easier to sell internally if it supports the corporate MIS strategy. If the technical product assessment indicates it would be appropriate to go “against the grain”, it is important to be aware that you are doing it.

The criteria for assessing a product's technical capabilities can be divided into three categories: software, hardware and support.

Software Considerations

- Operating system compatibility—what computer(s) does the product run on, and is this compatible with existing systems?
- Single point-of-entry—Can the user update all files from one entry point? Can the user make all entries at one place (the workorder)?
- Data base file structure.
- Custom report writer for end users—this is important for three reasons. First, no two organizations are alike; second, each organization's needs will change over time; third, network management sys-

Eight Critical Steps in Evaluating and Implementing Network Management Systems

NMS Input Data	Call Allocation and Billing	Equipment and Inventory	Directory and Message Center	Cable Records	Network Management	Work Orders
Extension #	X	X	X	X		X
Name	X	X	X			X
Title		X	X			X
Department	X	X	X			X
Division	X	X	X			X
Class of Service	X	X			X	X
Special Billing Codes	X	X			X	X
Location		X	X	X	X	X
Type of Set		X		X		X
Bridges		X		X	X	X
Ancillary Equipment		X		X	X	X
Installation Date		X		X	X	X
Associated Extension		X			X	X
Billed Cost		X			X	X
Pair Requirement		X		X	X	X
Call Pick-Up Group		X				X
Call Forwarding		X	X			X
Hunting		X			X	X
Node/Port/Card		X		X	X	X
Main Key Station		X				X
Coverage		X	X			X
Department Budget Number		X				X
Asset ID Tag Number		X				X
Cable Type				X	X	X
MDF/IDF/BDF		X		X	X	X
Outside Plant/Tie Cables				X	X	X
Jack Type/Pin Number		X		X	X	X
Trunk ID	X			X	X	X
Route Pattern	X				X	X
Time of Day	X				X	
NPA/NNX	X				X	X
Data Routing					X	X
Time/Date Reported						X
Time/Date Dispatched						X
Time/Date Technician On-Site						X
Time/Date Cleared						X
Technician ID						X
Trouble Found						X
Work Performed						X
Network Equipment					X	X
Circuit Number				X	X	X
Circuit Type				X	X	X
Circuit Routing (to/from)					X	X

Table 1. Sample data matrix showing network management system (NMS) input data on a module-by-module basis.

Eight Critical Steps in Evaluating and Implementing Network Management Systems

tems store vast quantities of information, and it is wise to take advantage of it.

- The number of product releases—it's best to obtain a system that has been "shaken down".
- Ease/frequency of enhancements—your telecommunications world will change, and the system must be able to change as well.
- Ease of interface with other software systems—existing as well as planned (LU6.2, financial systems, EDI, SMDR, etc.).
- Password security—there will be many users, and access should be controllable.

Hardware Considerations

- Disk space requirements—include room for growth over three years and then add another 33 percent.
- Number of terminals—include CRTs and printers. For example, estimate the number of persons processing workorders, trouble tickets, directory lookup, and bill reconciliation, as well as those performing network control and management inquiry. Again, plan out three years and then add 33 percent. This number, like disk space, is always greater than you think.
- Main memory—the amount required to run programs and still achieve the desired response time.
- Printer speed—keep in mind that call accounting is a very big report.
- Hardware interface to other systems—including polling system, trunk monitoring and/or mainframe.

Support Resources

Determine where you will derive technical support for the following areas once the system is up and running:

- Hardware.
- Software applications.
- Operating system software.

It is very important to understand *how* you will resolve problems in these areas. Many postinstallation difficulties are hard to diagnose between these three categories.

SUMMARY OF ACCOMPLISHMENTS: STEPS ONE, TWO AND THREE

At this juncture in your analysis you should have accomplished the following:

- Surveyed key departments to establish areas of concern, goals and management support.
- Analyzed the current methods of operation in each of these areas and highlighted both current problems and costs as well as future problems and costs in a three-year horizon.
- Identified the specific application areas of a network management system that will address the above identified concerns.
- Analyzed the information requirements to address these problem areas.
- Performed some general system sizing to determine processing categories to consider (i.e., micro, mini, mainframe, timeshare).
- Evaluated vendor's software to determine compatibility with information needs and its availability through a system.
- Determined ease and flexibility in use of software.
- Analyzed internal computing strategy and hardware/operating system support availability.

After analyzing these aspects, it should be possible to narrow the field of vendors to five or less. At this point, the user has determined which of these vendor's products *can work*, not that they *will work*.

The experience of both clients and vendors in the industry demonstrates that implementation is at least 50 percent of the battle. Thus, it is prudent to first narrow the search to a small number of vendor products that meet the organization's information and computing requirements. The final selection should be based on the vendor's ability to support the user not only during implementation, but on an ongoing basis as well.

Eight Critical Steps in Evaluating and Implementing Network Management Systems

STEP FOUR—GENERAL VENDOR ANALYSIS

Vendor analysis represents the beginning of another critical phase of the selection process. Successfully implementing a network management system requires more than just a product. It requires expertise in multiple disciplines, the availability of staff resources, and the experience to deal with the many surprises that develop both during implementation and thereafter.

When using these criteria to evaluate prospective vendors, it is advisable to involve key individuals that will be involved in the implementation process. This will help ensure supportive working relationships between internal and external personnel. The criteria for analyzing general vendor capabilities include:

- The vendor's years in the *network management software business*.
- The vendor's experience in your industry, with your PBX type, with your size system, and in your type of MIS environment.
- The vendor's staff—are they experienced enough in voice and data communications to understand your needs? Do they have sufficient software expertise, both in applications and operating systems? Do they have sufficient hardware expertise? This is critical to configuration design as well as ongoing problem analysis and resolution.
- The vendor's ongoing staff resource commitment—evaluate this by interviewing the primary support person assigned to your account. Analyze his/her ability as a project manager, availability in terms of time commitment, geographical proximity, and ongoing availability after implementation.
- Test and backup resources—evaluate the vendor's commitment to test software before shipment, resolve bugs, perform new development, and serve as backup to your department for polling, restoring files, etc. (Your department will depend on this.)

STEP FIVE—SYSTEM IMPLEMENTATION

The key to system implementation is to create an achievable plan of events, given the staff resources available. Implementing a network management system involves many steps and requires a lot of staff time. Most telecom departments are understaffed, and can ill-afford to take on an *extra load* for a 3-6 month implementation period. The answer to this

dilemma will be different for each organization, and it is an issue to resolve before beginning. Once again, the vendor can be a major asset (or liability) in this effort. The Example Case Study at the end of this report provides insights into some specifics to consider during implementation.

HELPFUL HINTS

If you are implementing a new PBX, train your staff on the system *before* installing the PBX and cable. During preinstallation, staff personnel can load data into the system while becoming familiar with the new equipment. Implementation can then be accomplished in an orderly fashion and will prevent the inevitable chaos of attempting to both install and load data simultaneously.

Initial implementation and training takes 3-6 months. Follow-up requires another three months. This cannot be accomplished in two weeks of training; moreover, staff personnel will never absorb it.

STEP SIX—DATA COLLECTION AND INPUT

Data collection and input are two major tasks when implementing a network management system. In many cases, thousands of pieces of information must be gathered from different sources and quickly entered into the system before the data is obsolete, since the telecommunications environment is constantly changing. It is critical to ask potential vendors how they plan to support this effort. **THIS CAN BE A MAJOR POTENTIAL PROBLEM AREA.**

FINDING THE DATA

First, determine what data is required to fulfill the requirements established during Product Assessment: Applications (Step Two). Review the data matrix developed during Step Two. Second, determine where this data currently resides. This often requires visiting various departments and "systems" within the organization. Many times there is a piece here and a piece there.

GETTING DATA INTO THE PROPER FORM

Once the data is located, take note of the form it is currently in, i.e., computer readable, manual records, no records, etc. Most organizations have all of the above. Next, determine exactly what form the computer readable data (software) must be in for entry into the system. To understand this requires close

Eight Critical Steps in Evaluating and Implementing Network Management Systems

work with the vendor. Determine how much data entry can be automated, and what must remain manual entry. Evaluate what resources it will take to get there, and finally, who (the vendor, the user, or both) is responsible for the varying tasks associated with automated or manual data entry.

The effectiveness with which this phase is performed will directly impact the *quality* of information produced. The axiom GIGO (garbage in-garbage out) applies here. The data collection and input step is the one typically most underestimated by clients, and the one which requires the greatest amount of planning and time.

STEP SEVEN—STAFF CONSIDERATIONS

Introducing a comprehensive network management system into an organization significantly changes the job descriptions of many individuals. The organizational structure and the philosophy of the entire department may undergo significant changes as well. This will, of course, vary from organization to organization. There are several areas to consider when assessing the impact of these changes.

First, what are the implications regarding the merging of voice and data today or in the future (or even sharing the same system)? If this is a possibility, then a system should be selected that will support both.

Second, does the implementation of a network management system result in the centralization of previously decentralized tasks? Are there political implications to this? Is there a way to phase in this change? If so, what are the hardware and software considerations in phasing in this change?

A network management system has a chicken and egg relationship to centralization. You sometimes don't know which came first but you always know the end result. This does not mean that all tasks need be centralized. Some tasks, such as workorder and trouble reporting, may remain decentralized, while the information management function becomes centralized. It is important to resolve this before system selection to ensure that the hardware and software will support the desired configuration.

Third, how will the implementation of a system impact the job description of each individual in the department? Are the individuals suited for these changes, and how will they react to these changes? What new job functions are created? Experience has shown that staff levels do not decrease during the first year after implementation. After year one, however, the staff grows at a slower rate than during the

“prenetwork management system” era. Furthermore, the existing staff is significantly more productive in terms of both throughput and effectiveness after the first year. Also, tasks tend to become more analytical in nature, eliminating the need for many redundant clerical functions.

SUMMARIZING STEPS FOUR, FIVE, SIX, AND SEVEN

At this point, you are in a position to select a vendor—or to confirm with the vendor that you have been working closely with that you have a plan that works. This plan should include:

- The product requirements that will provide the information necessary in an easy-to-use fashion to meet your goals.
- A comprehensive implementation plan which identifies all tasks, responsibilities, and time frames.
- Evaluation and resolution of all internal resource issues, pinpointing any necessary changes.

Now all you have to do is get it funded.

STEP EIGHT—COST JUSTIFICATION

The payback for successfully implementing network management systems can be significant. To determine the payback, however, the user must first assess the costs.

COSTS

Generally, costs fall into four categories: hardware, software, onetime implementation charges, and ongoing support costs.

Hardware costs are easy to assess for standalone systems. The costs for systems running on a mainframe are more difficult to determine, however. If possible, ask the vendor for quotes on both types. The standalone or departmental system will be a more controllable, onetime capital expense. Using a mainframe will generate a higher ongoing expense. Often, the choice between standalone/departmental system or following a mainframe approach is dictated by how the user's top management views each method.

Software costs include the total, actual cost of acquisition from the vendor. This may be a onetime license purchase price, or a monthly license fee. There may be additional fees for customization.

Eight Critical Steps in Evaluating and Implementing Network Management Systems

FINANCIAL	HARD SAVINGS	SOFT SAVINGS
Cost Control Usage (call accounting/network design) Equipment (inventory) Cost Allocation/Budget Bill Reconciliation	15-20% (of usage) 3-5% (of equipment) \$ associated with manual effort 7-10% (of equipment workorder charges, maintenance, trunk charges)	— — CFO must assign this value
OPERATIONAL Improve service to users/from vendors —Workorder/Trouble —Cable Management —Directory —Network Monitoring	— \$ reduction in labor \$ associated with manual effort 5% rebates for outages, reduce # circuits	Senior management must assign a value —More responsive to change and problems —Accurate information to customer/staff —Higher network availability
SYSTEM PLANNING Network/System Information	—	Senior management must assign a value to these —More effective capital expenditure decisions, planning —Improved ability to react to organizational change —Improve ability to identify strategic opportunities

Table 2. Estimated cost savings from using a network management system (NMS).

Implementation costs include vendor charges, overtime for staff, and any additional outside resources required. It is best to closely evaluate your implementation requirements. Often, the costs are underestimated.

Ongoing support includes software maintenance (which may be included in the purchase price or monthly license fee, at least for one year), hardware maintenance, and vendor support. Most vendors provide a toll-free hotline service; however, it is important to find out exactly how much advice is free, and when the vendor will start charging.

PAYBACK

Network management system objectives are very basic and are staples of effective management throughout all corporations. These objectives include *financial management*, *operations management*, and *systems planning*.

FINANCIAL MANAGEMENT

Financial management consists of cost control, cost allocation/budgeting, and bill reconciliation. By improving cost control through call accounting and network design, a network management system can

Eight Critical Steps in Evaluating and Implementing Network Management Systems

generate a 15 to 20 percent savings in usage. By maintaining accurate equipment inventories, a network management system may generate a 3 to 5 percent savings in equipment costs. Automated cost allocation and budgeting will generate a savings that is proportional to the previous costs associated with doing it manually. Bill reconciliation will generate a savings of about 7 to 10 percent of the total monthly equipment workorder charges, maintenance, and trunk charges.

OPERATIONS MANAGEMENT

Operations management consists of improving the quality of service, both to users and from carrier vendors. An effective network management system will improve cable management and workorder/trouble ticket processes so as to reduce labor costs. Automating directory services will generate savings in

proportion to the costs associated with manual processes. Network monitoring in itself should generate a savings of 5 percent from rebates for outages as well, making it possible to reduce the number of circuits needed.

SYSTEMS PLANNING

Systems planning assists the user in assessing the impact of organization changes on voice and data communications systems. It can also help to assess the impact of technological changes on the organization—highlighting both cost reductions and strategic opportunities.

Table 2 summarizes both the hard savings and soft savings associated with each of these three management areas.

EXAMPLE CASE STUDY: A NEEDS ANALYSIS DEVELOPED BY THE TELECOM STAFF OF A LARGE CORPORATION

The Telecommunications Department of Corporation XYZ is responsible for providing rapid, reliable, accurate and cost-effective voice communication service to employees. It is organizationally a part of data processing and, in turn, MIS.

The department's span of control consists of 15 private branch exchanges (PBXs) that service 30 buildings. There are more than 300 tie line circuits connecting the PBXs and approximately one thousand off-network circuits providing in/out communications to the PBXs.

The Telecommunications Department's equipment and circuits must satisfy the following requirements for its "customers" in the Home Office:

- To dial a telephone and be quickly connected.
- To get a clear circuit without unnecessary noise and errors.
- To have troubles quickly cleared.
- To have installations, changes and upgrades executed rapidly and smoothly.
- To be billed for services correctly.

- To be provided with accurate directory information.
- To be courteously and accurately handled by the operator.

To satisfy these requirements, the department established the following strategic policy:

- Provide pro-active voice communication network capacity to meet anticipated demands.
- Provide enhanced voice communications capabilities for customer service improvement.
- Implement internal programs to perform the span of responsibility.

Telecommunications is divided into five areas of responsibility within the organization: configuration management, performance management, problem management, change management, and administrative management. These areas reflect the functions that must be performed to satisfy customer requirements.

CONFIGURATION MANAGEMENT

Configuration management consists of identifying and controlling the hardware and software inventory. Hardware inventory includes all equipment (PBX, phones, etc.) and all circuits (tie line, WATS, CO, OPX, etc.). For both equipment and circuits, records must be kept on specifications,

Eight Critical Steps in Evaluating and Implementing Network Management Systems

locations, vendors, value and quantity. Software inventory includes numbering schemes, routing tables, feature tables and similar internal PBX tables (see Table 2).

PERFORMANCE MANAGEMENT

Performance management consists of diagnostic and fault monitoring, remote polling of the PBX for activity analysis, preventative maintenance and traffic analysis. The Telecommunications Department gears its performance management activities toward finding problems before they become apparent to the customer.

PROBLEM MANAGEMENT

Problem management means identifying and repairing failures, whether of equipment or circuits. Usually, failures are reported by the customer that is directly affected. However, the perceived failure may be only one component of the underlying problem. Problem management activities consist of reporting the problem to the appropriate vendor, scheduling and coordinating a technician, and providing for corrective maintenance.

CHANGE MANAGEMENT

Change management covers adds, moves, and deletes of equipment and services. The activity is primarily initiated by a customer's request for a change. It can also be initiated as a result of capacity planning analysis that indicates that a change is necessary. The initial step is to create a service order that identifies the hardware and/or software changes to be performed. The other activities consist of technician scheduling, order tracking, order modification, invoice reconciliation and payment, and activity reporting. The completion of the service order generates input to the configuration management activity and to the administrative management activity.

ADMINISTRATIVE MANAGEMENT

Administrative management consists of telephone vendor invoice payment, call accounting, usage and service chargeback billing, cost/performance analysis, directory (the department on-line and hard copy) maintenance and issuance, expense and capital budgeting and forecasting, training, and management reporting.

Call Accounting

Each month, the Telecommunications Department receives over 160 telephone vendor invoices, representing in excess of \$600,000. They must be approved for payment, the costs matched against the call detail records obtained from the PBX, and the resulting spread of expenses charged back to the appropriate customer.

Directory Services

Up-to-date telephone numbers for employees, departments and locations are necessary for the telephone operators to effectively and quickly transfer callers. The Telecommunications Department provides this information via on-line terminal access. The same information, in hard copy format, is provided on a regular schedule to company employees.

Miscellaneous

The Telecommunications Department continually analyzes the volume and demographics of on and off network traffic to determine the current amount of trunking available. The level of service performance provided is compared to the cost of providing the service to ensure an optimum mix is obtained. The department must produce control reports that provide management with the information necessary to ensure that the overall performance is positive.

The Telecommunications Department develops its budget by tracking actual expenses and creating appropriate forecasts.

Other responsibilities include providing for employees who require training in the use and capacities of the phone system. Additionally, the department must support a telephone console attendants service, which consists of a centralized facility where all outside calls are answered by operators located in the main plant. Over 650 calls per day are handled by each of seven console attendants.

AUTOMATION REQUIREMENTS

Managing telecommunications used to be a lot easier. Prior to the divestiture of AT&T and the department's subsequent acquisition of the Rolm CBX, the department consisted of eight telephone operators, two clerks and a supervisor. They depended on AT&T and the local BOC telephone to

Eight Critical Steps in Evaluating and Implementing Network Management Systems

add a circuit, move a phone or change Centrex extensions. There were fewer invoices to pay, there was no chargeback of costs, and no managed responsibility for the control or change of equipment.

Today the department consists of 25 employees, almost equally divided between exempt and non-exempt staff. The operators and their supervisor remain, but the control staff has expanded to address the five areas of telecommunications responsibility discussed previously. These areas have appeared and matured since the initial Rolm installation. The scope of responsibility, amount of data processed, value of equipment and services, and pace of activity has equally increased.

These increases have been addressed by various measures, most notably staff increases. Areas that were not previously managed internally now require staffing. At the same time, the department introduced several limited automated solutions to handle the increased volume of information. Piece-meal computer approaches have provided stopgap methods which created as many problems as they have resolved. Unfortunately, the bulk of the information and entire areas of responsibility continue to rely on manual methods.

On the following pages, each of the five areas of responsibility are described in detail to identify the shortcomings of the current methods of operation.

Configuration Management

A system's configuration is a map of all the hardware within a switch and an outline of all software controlling the functions of that switch. The software map consists of such information as station characteristics, class of service, feature tables, routing guides, hunt groups, ACD groups, and trunk display groups. In addition, the software also includes the system parameters that operate that particular switch i.e., when lights flash on a phone, how long a call is on hold, etc.

A switch is configured for an expected number of stations and circuits calculated for that particular building. When there is an increase that exceeds those expectations, a reconfiguration is necessary to change the software parameters to accommodate the additional numbers. This reconfiguration process is both time-consuming and very costly. Because of the time and cost involved, the Telecommunications Department attempts to bundle as many system changes as possible into one reconfiguration.

If the requirement is for a software change only, that change is made online into the system and becomes the new software record immediately. Written documentation on all these software changes must then be forwarded to Rolm support center, since a record of every change in a switch must be kept on hand in case of a major failure. (Presently this documentation is handwritten and mailed to the service center.) Failure to send this information can result in a service contract dispute.

The change to the configuration software is a separate function of completing a change. However, the paperwork for accomplishing this task is the same workorder as completing the physical move of the equipment. A handwritten document or a verbal order to Rolm authorizes the change (the procedure is not yet automated). Like the move and change function itself, it is time-consuming to administer and very difficult to reconcile.

The continual movement of employees from one location to another, each with varying requirements, has a significant impact on the switch capacity at each location. In evaluating possible network management systems the following questions must be addressed:

- Are there sufficient boards available with the equipment?
- Are there available direct inward dial numbers to assign to customers?
- Are there a sufficient number of tie and other trunks to meet their calling requirements?

These questions must be addressed from the start, since overlooking even one of these items can result in either substantial delays in service or deteriorated service.

Performance Measurement

The Telecommunications Department performs a number of analyses to monitor the performance of telecommunications services elements, both from a cost and response perspective.

The department measures traffic volumes by utilizing several different sources of information. Specially developed polling programs, running on PCs, collect and analyze sample traffic statistics from the PBXs on a weekly basis. These programs are useful in determining the capacity margins of groups of circuits between locations and to/from

Eight Critical Steps in Evaluating and Implementing Network Management Systems

the telephone companies, but provide no information on individual circuits. From these programs, the department can determine whether there are enough circuits to handle the traffic—however, the programs do not show when individual circuits are malfunctioning. Another measure of traffic information is available from the department's call accounting systems, which offer some summary reports but do not provide sufficient information about circuit performance.

Another custom program collects statistics on console operator volumes and response times. The system is limited to the previous day's activity, and does not give information on trends or larger periods.

The Telecommunications Department also maintains an analysis of long distance costs and volumes which determines how completely the department charges back those facilities. The same analysis is used to monitor the overall cost per minute trends. As with the other analyses mentioned, this one is not directly connected with the billing system.

There is no mechanism or analysis used within the department to measure internal problem resolution response time nor that of the vendors.

Problem Management

Any individual within the company may contact telecommunications to report problems either with the telephone itself or with placing calls. Upon receiving a repair call, the information needed to obtain for proper problem determination includes:

- The name of the user experiencing the problem.
- The user's extension and building.
- The user's location within the building.
- The type of telephone.
- Complete description of the problem.

Information obtained from these initial contacts helps to identify whether the problem is an isolated station problem, software problem, or a switch or facility problem. This identification process is extremely important since precious time can be saved by accurately diagnosing the cause.

From the above information, the department determines whether to report the problem to Rolm,

to another vendor such as AT&T, or whether it can be resolved with internal personnel. (All trunking problems with the exception of a few unique services are reported by Rolm to the vendor).

What's wrong with this process? It is difficult to retrieve a listing of current problems. Whenever someone calls to question a repair, department personnel must look through the handwritten repair log to find the repair. Even when the report ticket is found, the status of the problem isn't available.

Although one person is responsible for reporting and tracking repair problems, everyone in the department is responsible for answering repair calls if that individual is unavailable. Because there isn't an easy way of inputting or retrieving information, some repairs may be overlooked or multiple calls for the same problem may be made.

Tracking repairs manually results in a lack of historical data and a lack of station information. This generates two serious problems. First, in the event of recurring switch or trunk problems, it would be desirable to discuss historical data with the vendor in order to reach final problem resolution. Second, department personnel could refer to historical data and station information when customers call in to report problems. Because information is not readily available about the individuals calling in a problem, the department cannot review the specifics of their equipment while they are still on the telephone. Problems which are software-related or related to an individual moving a telephone without telecommunications authorization are reported to Rolm. In prior years, this wasn't a billed item, but that is no longer the case. Because of the significant number of customer-caused problems, Rolm now charges the department for time involved in researching these problems.

Change Management

Information regarding the addition of new equipment, the move or change of existing equipment, software changes, or a combination of all of these is relayed to Telecommunications verbally or by written request. The information required to prepare the work order includes:

- Who is moving.
- When the move will occur.
- Cost center.

Eight Critical Steps in Evaluating and Implementing Network Management Systems

- The building.
- Software changes.
- Hardware requirements.
- Extension numbers.
- New location.
- Old location.
- What exists today.

All necessary information is handwritten on a form and either called in to Rolm or given to department personnel for completion.

At the time of assigning the workorder to a particular group, a completion date is given to telecommunications. This date is then verbally relayed to the customers requesting the move along with any personnel in the facilities department who need it for the completion of their work. After the work is completed, Rolm signs off on their order and subsequently that paperwork is used to verify the final invoice.

The entire process is manual. Complete information must be gathered about the existing station configuration as well as what is proposed before a workorder can be submitted for action. Because of the difficulty of retrieving data, workorders often contain only the most vital station information. In many cases, even the names of all the customers moving are not included because it is too cumbersome. Retrieving vital information such as cost center or locations from other automated files within the department (i.e. chargeback system, directory, etc.) is not a viable option since the means for retrieving this information is unavailable to the configuration specialist. While this isn't necessarily a problem for the processing of the workorder, it does mean that information does not get passed to the other administrative groups within the department. The result of this inability to easily pass accurate data is that the Telecommunications Department does not have correct telephone numbers or locations, and cannot charge users correctly for their equipment or long distance calls. Also, the department loses control of the location of its telephone equipment.

Administrative Management

Call Accounting

Currently, the department operates two PC-based call accounting systems for recording and processing long distance calls placed from the corporate network. These systems are connected by cable to the two PBXs at the main plant and headquarters. These are the locations that provide outgoing long distance service to the entire network. When a call is processed by the switch in either location, a record of the call is sent to the PC to be stored and eventually costed and charged to the appropriate cost center. The call accounting systems operate as they were intended, but they cannot provide all the functionality needed to properly control our long distance and other calling expenses.

The weaknesses of the present call accounting system are primarily a function of its relative simplicity and the relatively limited storage and processing power of even an IBM PC/AT. The PC systems are not capable of adequately handling the volumes of calls that should be processed to properly manage current usage.

For instance, the elimination of authorization codes (a project to be coincident with the automation system implementation) will increase chargeable long distance calls by 50 percent. This is because the use of auth-codes doesn't permit the department to charge back ADN (Automatically Dialed Number) calls. The department must compensate by over pricing all chargeable calls, which results in an inequitable distribution of costs.

The call accounting system requires that a directory of auth-codes, names, and departments be maintained separately from the online and hard copy directory systems with the inevitable inconsistency and triplication of some data.

The present system does not track changes in station or auth-code assignment by date. Updates to the current month's records result in the loss of accurate historical information. Changes that take place in the middle of the month must be charged to an entire month, or handled through some more complicated billing mechanism. The best way to manage this without a software change would be through the maintenance of numerous archival copies of our databases.

The costing methods available on the existing system are not flexible enough to support equitable charging of long distance calling. The system has

Eight Critical Steps in Evaluating and Implementing Network Management Systems

reached a capacity limit on the number of route costing parameters that can be used to cost calls differently. The department is also unable to cost calls at anything other than daytime rates. Calls placed at off-peak periods (before 8 a.m., after 5 p.m. and on weekends) receive no discount, and hence, no incentive.

The present system does not provide the necessary features to charge back all telecommunications expenses, so a separate billing system had to be constructed to consolidate the half dozen expense elements. As a result, the interface between the call accounting systems and the billing system is a clumsy patchwork of diskette transfers, custom programs and manual intervention.

The existing system(s) cannot accommodate the direct input of calling card, international and operator assisted charges from tapes provided by AT&T and the local BOC. Instead, summary totals are manually entered into the billing system, and details are available only by making copies of paper telephone bills.

A number of useful reports are not available directly from the call accounting system. Most important of these is the lack of a standard report of the most frequently called numbers.

Chargeback Billing

Each month, the Telecommunications Department produces a Journal Voucher (JV) that charges over 500 cost centers for about \$100,000 of expenses from numerous sources. In 1985, this largely manual operation was replaced by a customized PC database application that shares hardware with one of the call accounting systems. It combines chargeback data for equipment allocations, long distance calling, calling card charges and pass-through expenses into a single total for each cost center, and generates a JV printout. It also provides some detail and summary reports that are collated and mailed to the appropriate cost center managers for review.

The system suffers from some of the same weaknesses as the call accounting system. Primary among these is the inability to date-stamp changes to equipment records. As with call accounting, the reassignment of a phone can only take place at the end of a month for billing purposes, and the historical assignment is lost. Historical information isn't available with the present system, and the sheer volume of reports that it produces makes it impractical to maintain hard copy back-ups. Once

reports are sent to cost center managers, additional copies are available only until the next billing is due, usually a matter of a week or so after a mailing.

The billing system shares information with the hard copy employee directory, but is not connected with the auth-code or operators' on-line directory. A new employee's phone can result in the updating of four or more separate files on three different computers. Two of those files determine how their telecom charges are billed.

Because the billing system's equipment database does not interface with telephone order processing, a phone installation, removal or reassignment is frequently not reflected in the database. This means that the associated cost center is not properly billed if the person in charge of a given database is not informed of a change in a timely manner. Consequently, the department has hundreds of phones which have no cost center assignment, and perhaps hundreds more that are incorrectly assigned. Those phones that are identifiable are therefore charged more to cover those that aren't.

One major task for the department each month is the collation of the 500 cost center billing summaries. Long distance details from the call accounting systems comprise two of the six stacks of paper that must be married together manually. The others are a summary page, extension, call card and pass-through listings where applicable. A fully integrated system could eliminate this chore.

Other Activities

Among the miscellaneous services provided by the department is the occasional request from security or employee relations to provide information on special call related problems. This frequently takes the form of monitoring call records for evidence of obscene or abusive phone calls. Unfortunately, the tools available to the department are not always adequate to fulfill the request, usually at remote locations, where local call detail is not available.

Directory

The Directory functions provide accurate telephone number and location information on 5,000 home office employees, 800 home office departments, and 200 subsidiary and field organizations to the console attendants, individual employees and the expense chargeback function. On the aver-

Eight Critical Steps in Evaluating and Implementing Network Management Systems

age, it processes 450 changes to these records each month. The console attendants must have subsecond real-time access to the information and a hard copy should be provided on a quarterly basis.

Change information is obtained in a variety of ways. PDS reports, copies of PCA forms, change request forms provided to customers, interoffice memos, telephone calls and even personal visits are all methods that are used to obtain current information on employees, departments and organizations. At the same time, some information is obtained from the input and/or results of the change management and chargeback billing functions.

The information is initially stored in a manila folder. The folder is labeled as to the type of information it contains and the computer system and file that will be updated with the information. After the file is updated, the folder is marked "completed" and discarded.

There are two distinct and separate directory systems and two additional related files within the chargeback billing function that contain employee, department, and organization information.

The data from the change information forms is entered into the appropriate file(s) in both of the computer systems and the related chargeback billing files. In both systems, the date of update is entered for identification purposes only. The information changed in the online system is immediately available to console attendants.

The changes to the hard copy directory can only be made during the window in the chargeback billing system between the completion of the previous month's detail report preparation and the start of the current month charge determination. It takes about four to six weeks to create the hard copy directory. The departments notifies all employees that it will soon issue a new directory and states the cut-off date for accepting changes. A significantly larger amount of change notification forms

are received that represent previously unannounced changes, such as married names, swapping of offices and extensions, and new employees.

After the updates are made to the four files within the hard copy system, the files are electronically transmitted to the graphic arts department. Proofs are returned for verification and, after final corrections are made, the plates are given to In-house Printing where 7,000 copies are printed.

The Directory process is at best, cumbersome and at worst, redundant and terribly inefficient. With the exception of new hires, terminations and name changes, the information on extension and location changes is available in the change management function. Unfortunately, it is usually inaccessible due to the sheer volume of workorders and because it is stored manually. Much effort to identify directory changes can be avoided with an automated feed on an occurred basis using an effective date. The manual storage of forms, the logging of the form's status, and the resulting update proofing of the multiple files would be unnecessary.

A single system with interrelated files would alleviate the multiple putting of the same information upwards of six times. Manual input to the chargeback function would be unnecessary as the directory file could serve as the source of input and avoid a separate file. Correspondingly, an automated system would avoid repetitive verification which now ensures the common use of information throughout the multiple files.

The capability to update as changes occur using an effective date would alleviate the situation where, through employee absence or work on higher priority assignments, changes are not processed and incorrect information is retained. By eliminating the need to update and verify file accuracy, console attendants would have access to immediate and accurate information. An automated system would also greatly reduce the time necessary to prepare and issue a hard copy directory.

Designing Network Control Centers for Greater Productivity

This report will help you to:

- Plan for the implementation of a network control center.
 - Evaluate the control center's ergonomics.
 - Anticipate future control center needs.
-
-

Despite the fact that networks play a key part in corporate activity, little attention is being paid to the nerve center of this function: the network control center. The corporate approach to control center design today runs the gamut from models of efficiency to routinely backlogged, poorly utilized resources. Although a properly designed environment is not the final answer to maximizing human and hardware productivity, recent studies have confirmed that it can go a long way toward improving a less-than-ideal situation.

Part of the problem is that private networks are still novel at many companies, so the concept of a specialized place dedicated to network operations tends to get short shrift. Another factor is that there are few historical studies of optimal command center design.

At the heart of the problem, though, is the intensity of the human/machine interface. The problem is not without precedent. Consider recent technological history: The same problem has already been faced in the scientific and financial communities, both of which provide an excellent starting point for deciding how to build a better data/voice network control center.

Perhaps one of the most vivid images of the last two decades is the rows of NASA controllers monitoring individual data terminals during a spacecraft launch. With at least one monitor per person and numerous

monitors per spacecraft function, there was need for a well-organized environment that allowed both individual monitor attention and attention to the project's full progression. This was accomplished by arranging low banks of monitors, one per analyst, facing wall displays providing an overview of the entire project.

If the ergonomics of control center design seems to be overstatement of the obvious, consider the following specific design parameters:

- What is the optimum height of the monitor so that it can be the primary focus but still be looked over if needed?
- How deep should the counter in front of it be?
- How close can the monitors be to one another?
- How close can the people be?
- What accommodations should there be for telephones, lighting, and wire management?
- How do you provide for equipment service and changes?

All of these issues contribute to formulating preliminary mechanical-design needs, but equally important is an understanding of each job function and specific needs associated with carrying out that function. To illustrate the importance of function, consider a well-designed trading module used by financial institutions.

This Datapro report is based on "Designing network control centers for greater productivity," by Kerry Kosak, Descience Corporation, from *Data Communications*, April 1988. © 1988, McGraw-Hill, Inc. Reprinted by permission.

Designing Network Control Centers for Greater Productivity

The job of trading requires quick and constant access to two primary sets of equipment, the on-line monitors and the telephone. Further, because traders often work in groups, a systematic means of clustering people in either an open floor plan or in rows is a prerequisite, and equipment changes are constant.

Therefore, a well-designed trading module requires flexibility for differing floor plans, modularity to handle diverse and changing equipment, a high level of equipment density, and adequate space for documentation and personal storage.

To add a little more sophistication to the design equation, factor in the trend toward using the trading floor as a showplace in many financial institutions. Suddenly, aesthetics become relatively important and hidden-wire management, as well as data, telephone-line, and power access must all be juggled without jeopardizing crucial mechanical and ergonomic design elements.

ORGANIZATION

Designers of specialized environments, such as those seen in financial institutions, have amassed a body of experience that other on-line applications can use. When the lessons learned on the trading floor are combined with the ergonomic work done by many companies and institutions, a substantial pool of data is available that can be applied to the specific needs of the network control environment.

The first thing that is obvious when entering a network control center is the ratio of monitors to personnel. In the financial environment, there are generally three or four monitors per person in a defined workstation. In the network control center, it is more common to find several people interacting with common monitors.

This arrangement is further complicated by the fact that different staff operations often share the same space. For example, it is not uncommon to see systems analysts, technicians, and a help desk all in the control area, all sharing the same hardware to one degree or another. Therefore, organization, ease of presentation, and access to a high density of monitors become critical concerns—problems that can be easily complicated when space is restricted or at a premium.

The design solution is vertical stacking. However, since a person may view a stacked monitor from either a sitting or a standing position, optimizing the dimensions and layout of the stacks is critical. Setting the design dimensions compatible with people from

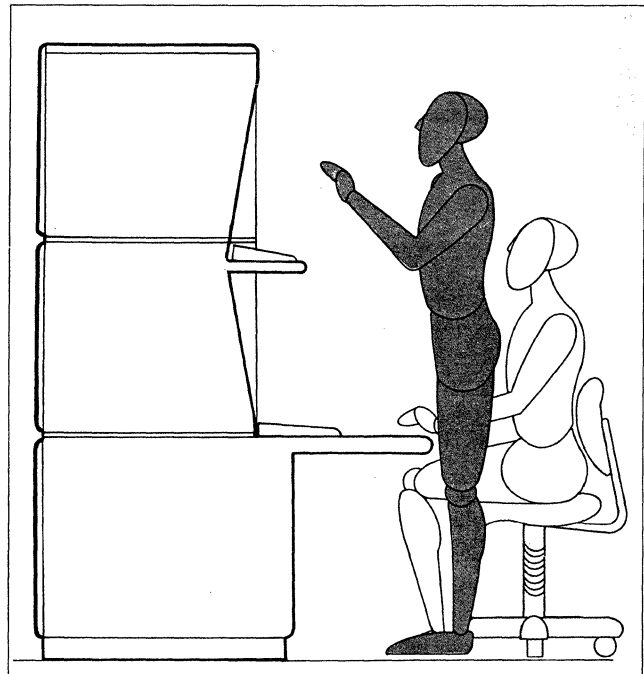


Figure 1. Viewing tiers of monitors can often lead to fatigue and neck strain, but when the screens are angled 10 degrees from the vertical, they become easy to see from either a sitting or a standing position.

five feet to six feet, two inches tall will cover approximately 95 percent of the population. Beginning with a standard desktop height of 29 inches supporting the first tier, the second monitor tier will top out at just about six feet. This height precludes realistic options for a third tier, since it would be almost impossible to access and difficult to see from a sitting position.

Designing the monitor location in the second tier for viewing from a seated position creates other problems. First, because every monitor generally requires that a keyboard be placed in front of it, there should always be a shelf with enough depth to handle the largest keyboard in use (approximately 12 inches deep). Therefore, the viewer will always be looking up and over the shelf. A straight vertical presentation of the second tier monitor can be easily viewed by someone in a standing position but is difficult to see when sitting.

The optimum relationship between an upper and a lower screen is to have both fall within a 20-degree cone of vision of the seated analyst. This requires that each monitor screen be angled approximately 10 degrees from the vertical, which facilitates seeing both screens from either position (Figure 1).

At the network control site for one of the largest insurance companies in the United States, these design techniques permitted the consolidation of three

Designing Network Control Centers for Greater Productivity

staff groups into a control area 40 percent smaller than the area previously reserved for the systems analysts. With the addition of two more groups to the same work area, management reported that the ability to quickly pull together people from diverse disciplines for any given problem resulted in a higher level of efficiency and improved customer service.

ACCENTUATING THE FUNCTIONAL

Beyond organizational and staffing aspects, the functional and mechanical needs of the control center hardware must be addressed. The mechanical layout must provide easy access to monitoring equipment in both tiers, and such things as data lines, power, and air-conditioning access are all essential aspects of the well-designed network operations center. The design of the screen-support stack should be adaptable enough to handle a diverse range of monitor sizes and types and be structurally solid enough to avoid any chance of tipping over when loaded with monitors. These are primary concerns in any design that is freestanding in an open control room space.

LIGHTING AND GLARE REDUCTION

The most overlooked aspect of control room design is lighting. Generally speaking, most control rooms have too much light and the wrong type. At the center of the problem is screen glare. The reflective nature of the monitor-screen surface in a highly lit environment can become a major distraction to an operator at a terminal. Part of the problem can be alleviated by using a nonreflective surface to cover the monitor, although high levels of light can still wash out the data.

Another aspect of the lighting problem is that despite the fact that it is easier to read a screen in a low-light environment, analysts have reading, writing, and keying tasks to perform on an ongoing basis. The optimum solution for these seemingly conflicting needs is not as complicated as one might imagine. The key is simply to provide controlled light only where needed. This can be accomplished by providing adequate task illumination on the horizontal work surface, supplemented with adjustable indirect ambient lighting kept at minimal levels.

The best solution appears to be built-in direct lighting from a position overhanging the work surface. By using louvered grilles, light can actually be channeled to spill on the horizontal plane of the desktop in front of a monitor. This will provide adequate task lighting for operators without creating a wash or reflection on

the monitor screen, while allowing ambient lighting to be reduced significantly.

What happens when area lighting is poorly planned? The experience of a large East Coast securities company sheds some light on this type of problem. The company's new command center was designed to be situated in the room with the mainframe and the support hardware. It was placed on a platform six inches above the computer floor. The additional height, combined with a generally high level of fluorescent lighting, immediately caused glare problems, even though nonreflective screens were in place.

Although the network-operations personnel were looking forward to an integrated control environment, the lighting conditions caused problems. Most of the personnel were frustrated, a few openly complained that they could not work there. The lighting problem added more stress to the settling in and bringing network operations on-line.

The lighting problem was eventually corrected by reducing the overhead light directly above the control environment. But it did require electricians to reopen ceiling access for a period of days, thereby delaying the start-up of a fully functioning installation.

In the grand scheme of things, glare may seem an oversight, but only because it was easily correctable. One only had to hear the complaints of the analysts to understand how critical the proper design of lighting is to the entire command center.

The lighting issue underscores the single overriding goal in control center design: Emphasize the data, not the hardware. In fact, in the perfect environment, all data is contrast-enhanced and well presented; all hardware is practically invisible and almost totally silent.

A well-designed network control room eliminates all visual distractions. Irregular monitor shapes, cabling, and other hardware are obscured by enclosing the monitors in a dark, nonreflective cabinet. However, enclosing monitors requires forethought to provide functional solutions that facilitate accessibility, ventilation, and wire management.

Another major distraction is noise. Noise is important when deciding on the proximity of the command center to printers and other distractions associated with network hardware. Many companies have used glass enclosures contiguous to the actual hardware area to set off the command center. This helps organize space as well as cut down the noise level in the command environment.

Designing Network Control Centers for Greater Productivity

To analysts with so many critical responsibilities, providing an area secure from hardware noise is a definite advantage. In this environment, enclosures can reduce the low-frequency hum of the monitors. There is also a variety of enclosure setups that can reduce printer noise by as much as 90 percent.

ADVANCE PLANNING: THE CRITICAL CRITERIA

How can networking professionals plan for the future in operations center design? In most instances, the impetus for designing a new command environment is the projected growth of the network. At the command center, growth is compounded by other concerns, such as the phasing in and out of computer hardware and telecommunications equipment. All of these factors have an obvious impact on proper design.

One way to get a handle on these issues is to examine an individual company and follow its planning and design decisions. For security reasons, the company will be called the ABC Corp. It is a holding company for firms in a variety of related industries.

Several years ago, following a change in management, ABC began an aggressive program of acquisition. Some companies remained self-contained subsidiaries, while others were folded into the parent company for economic reasons.

Under this regime, the company's revenues have increased more than fivefold, while management information systems (MIS) capabilities have increased proportionately. At the beginning of the acquisition process, the MIS function was performed in a single large room, but it soon became a series of installations spread throughout a midwestern city. Before MIS consolidation began, three mainframes were located in different parts of the city.

One of the first steps for integrating MIS functions was to build an advanced network consisting of both fiber optic and leased lines. The experience gained from this venture gave MIS operations personnel the confidence to assume responsibility for monitoring and maintaining telecommunications throughout the company.

Three years ago, ABC management realized that it was approaching critical mass and began the planning process for a centralized command environment that could match projected company growth into the next decade. Three overriding concerns directed their efforts:

- Accommodating all present and future mainframe needs.
- Consolidating the monitoring setups for telecommunications and MIS environmental and support services.
- Planning a three-phase expansion to facilitate growth.

ABC's new data center, only recently completed, is an impressive example of the importance of advance planning. The command center is the geographic center of an installation that takes up the entire floor of a medium-size office building. As such, the command center serves as a functional hub for all the diverse operations it monitors and maintains. The semicircular, double-tiered monitor banks built of modular enclosures, the back wall containing all security, power, and environmental monitoring equipment, and the nearby lounge area for the analysts are all the result of comprehensive planning.

Within the security area, taking up roughly half the floor, are all printing, mainframe, disk-drive, telecommunications, and power resources. The other side of the floor houses offices for the management and support staff. In each of the areas, color schemes, furniture, wall treatment, and lighting have been coordinated for continuity as well as for the needs of specific working environments.

Despite the striking visual impact, an equal amount of precise planning lies behind the walls to accommodate expansion. First, power, cooling, and environmental monitoring and control have all been overbuilt, allowing the raised floor area to be expanded by more than 20 percent with easy access to needed services. Non-weight-bearing modular walls have been used extensively to allow expansion and reconfiguration of the existing center. All of these elements are geared for anticipated MIS growth, which is expected to double again by the end of the decade.

ABC hired a consulting contractor to both plan and execute the work. According to the network-operations manager, this decision was made early in the discussion phase, since ABC realized the importance of support services and that control center design was one area where it lacked the proper expertise.

ABC knew its specific operational needs, but it relied upon the consultants to provide recommendations about power and environmental monitoring—

Designing Network Control Centers for Greater Productivity



Figure 2. As the network control center becomes equivalent to the nerve center of an entire company, it becomes the paradigm of corporate culture and capabilities. The layout of Burlington Northern's network operations center made monitoring and control operations easier, while the improved aesthetics made the command center a company showplace.

including water detection, main and backup air-conditioning, and Halon fire suppression—as well as security.

A majority of the support-function monitors has been integrated into the back wall of the center. This was done to make the command center a completely functional nerve center in the event of an emergency, an operational consideration often overlooked in the design phase of new centers. Such features as multi-zoned fire suppression, two-tiered security, and redundant cooling are added dimensions to traditional plans that can become important for both expansion and emergency contingencies.

HUMAN ELEMENTS

Hardware is not the only important issue, however. As the example of area lighting demonstrated, the needs of the personnel must also be given high priority. Once the explicit design criteria have been met, attention should be paid to the specific needs of the individuals.

The personnel staffing the center should have adequate storage for personal effects and basic supplies. Additionally, storage space for hardware and software documentation should be both plentiful and accessible.

Where monitor banks are stacked two tiers high on a desktop base, adequate work space should be accounted for in the design. However, desktops that are part of the monitor tiers should be set back deeply enough to accommodate either two keyboards (front and back) or a single keyboard and standard binders

of documentation. This allows the analyst to use the documentation and keyboard simultaneously on the work surface without requiring cumbersome under-counter pull-out trays or using the analyst's lap as an extension of the work surface. This need is apparent when a software analyst works for several hours without adequate surface area.

Consideration should also be given to where the analysts can spend their time when not specifically working in the control environment. In highly secure areas, one might follow the example of ABC, which built an analyst lounge right off the main secure corridor between the disk-drive operations area and the command center.

Operators' performance can also be enhanced by providing safety features such as structural stability, fire-proof construction, abatement of potential radiation, and rounded edges on all exposed surfaces.

THE CORPORATE IMAGE

As the network control center becomes equivalent to the nerve center of a company, it becomes the showplace of corporate capabilities. So, beyond space, mechanical, and ergonomic needs, aesthetics is playing an increasing role in the location and design of the command center. There is no question that an aesthetic environment can become an integral part of a company's identity, but like all other aspects of the optimum design, it must be factored into the design equation without detracting from other needs, as in the Burlington, N.H., network control center of Burlington Northern Railroad (Figure 2).

Designing Network Control Centers for Greater Productivity

Finally, under all circumstances, the command center must be designed to be redesigned. It must be able not only to account for an ever-changing array of monitoring and technical setups but also be adaptable to changing space configurations and provide hardware and service access. Telecommunications, power, and ventilation also must be considered, as should cable management.

For example, at one IBM manufacturing plant, after just 18 months, the increasing responsibilities of network operations forced a 40 percent expansion of the network control center. By using a modular design that allowed for easy expansion and reconfiguration,

the expansion could be made cost-effective and carried out in a fraction of the time required for an entire custom-built design.

In the final analysis, command center design is not simply a function of space allocation and furniture. An enormous amount of advance planning is required for centralized control centers, especially in departments that have significant growth projections. To be effective, questions of morale, hardware, peripherals, and even the building's architecture have to be considered. Eliminating any single aspect can substantially decrease the center's effectiveness and longevity. □

Homegrown Network Management

This report will help you to:

- Plan a network management system that is effective today and will accommodate future standards.
 - Evaluate currently available network management systems.
 - Develop strategies for implementing a hybrid network management system.
-
-

USERS CULTIVATING HYBRID METHODS TO MANAGE NETS

Faced with difficult conditions and a changing environment, farmers long ago developed hybrids to improve the yield and quality of their crops.

Today's network planners and users have similar problems. With equipment from a variety of vendors using different media and divergent protocols, users must develop effective hybrid network management systems to handle the often harsh communications climate.

The term "network management system" encompasses a variety of products and services. Recently, net management systems have evolved from a collection of modules, such as IBM's Network Communications Control Facility (NCCF) and Digital Equipment Corp.'s Network Control Program (NCP), to more comprehensive systems, such as IBM's NetView and DEC's Network Management Control Center. However, since these systems are managing only specific network elements, the term network element management systems (NEMS) will be used when describing them.

This Datapro report is based on "Homegrown Net Management," by Chuck Papageorgiou, United Parcel Service, Inc., from *Network World*, October 1988. © 1988 NW Publishing, Inc. Reprinted by permission.

LIMITATIONS

The limitations of these systems remain the same as in the original module offerings. Vendor-developed NEMs are designed primarily around the vendor's architecture with occasional hooks into other net management systems or other vendors' equipment. They concentrate mainly on managing elements such as modems, terminal controllers, and front-end processors.

These NEMs control and manage only specific areas and elements within the variety of networks that compose the corporate network.

Consequently, vendors often overlook the need of the user to manage all the different networks that form the whole without worrying about the protocols or architectures that it comprises. The user is left with incompatible NEMs that need, in most cases, separate staff, with different training and levels of expertise, and reside within vendor-specific subnetworks accessed through different network operator consoles.

PRESSURE FOR CONNECTIVITY

The pressure from large corporate users for complete network management forced the vendor community to form alliances and provide interfaces to other systems, in a way similar to the approach IBM takes with NetView/PC. However, these interfaces are

Homegrown Network Management

mainly patches, and their functionality is limited to alarm collections and other reporting functions.

The general consensus is that no single, commercially available net management system can perform the task of managing a multivendor, mixed-architecture network. It is therefore essential to view network management in a different light. Any planning or discussions that take place must consider the following:

- The user must acquire a number of NEMSSs in order to provide efficient management for all the network elements. As a result, the staffing levels and training requirements for the personnel in charge of managing the network increase dramatically.
- The user must assume overall responsibility for successful integration and must follow a careful NEMS module-selection process.
- Before the integration process begins, the user should take under consideration the imminent arrival of international standards for network management.
- As a last resort, the user must also consider implementing a single-vendor corporate network. This, however, is not practical since it limits the options available to the user.

A review of offerings from various companies shows that two design architectures appear to be most commonly used by vendors.

The first approach is based on the notion that a net management system should focus on path and network-availability control. This approach is taken by modem and other communications equipment vendors.

A system designed in this way allows the network operators to identify disrupted paths or congestion points in the network without needing to know which applications or users are affected.

The system provides alarms and status reports about the network components, such as lines, modems, multiplexers, and other equipment, in real time. The status of the entire physical network is available to the network operators at a glance using graphical representations.

This type of net management system becomes an extension of the physical network, and it ignores the network's reason for existence—serving the end user, whether an application or a person. This method is valuable for identifying and correcting physical net-

work problems, sometimes before the user even notices them. However, it does not help in cases where the problem is application specific, that is, in the software. Eventually, this net management system's inability to associate network paths with applications becomes a liability.

Most end users don't know or care what piece of the network they are connected to; they are only concerned about the applications they are using. From a user's perspective, if for some reason an application is not accessible, the network is considered unavailable, and the details do not matter.

A potential problem with this type of net management system is the proliferation of packet-based protocols and peer-to-peer connectionless communications. Point-to-point circuits are replaced by virtual paths, the network is used as an intelligent router, and sessions can be active even if paths and nodes fail. When that occurs, a node or link fails, and the net management system will identify it and produce an alarm, just as it is supposed to. Because it is concerned only with physical paths, it will not be able to provide the application with reconfiguration or fault management.

Net management systems of this type are hardware and modem control systems and other equipment-monitoring facilities, such as the Dataphone II Network Management System from AT&T.

SESSION-ORIENTED APPROACH

The second approach is to design a net management system that focuses on session control. This approach allows the network operators to identify and monitor calls to an application through the network without needing to know what route the call took in order to reach the application. A typical example is the functionality built into IBM's NCCF. By having access to the VTAM facility and NCP, NCCF permits a network operator or program to monitor and control sessions. When a session is lost, the network operator can re-establish it or, more specifically, reactivate it.

However, if the problem is a bad line that requires dial backup, the operator has to use a path-control net management system, such as the AT&T Dataphone II, or manually initiate the dial backup procedure. When the link is re-established, the operator can reactivate the session. However, NCCF will not be aware that the session has been rerouted to a dial backup facility because the path-control net management system is concerned only with the hardware and doesn't communicate with applications. In this case, the tool

Homegrown Network Management

becomes an extension of the logical network, and it, from a management point of view, ignores the physical network.

The session control approach provides a very good answer to the familiar cry, "My system is down." If the application is up and running, the net management system can direct the network operator to search for the problem somewhere in the transport or physical areas of the network. Since this type of net management system does not help identify where the call gets blocked or fails, individual physical network troubleshooting mechanisms have to be invoked.

PHYSICAL/LOGICAL INTEGRATION

Since both the aforementioned design approaches to network management stop short of a complete net management system in some way or another, a network management system designer must consider a mixed solution to the network management problem. By combining the physical management capabilities of the first approach with the logical capabilities of the second, a system designer can implement an efficient management system that controls the entire corporate network.

Some vendors are addressing the physical/logical net management system integration problem. Some even provide "complete" network management systems, both from the logical and physical points of view. IBM's NetView is a good example. It includes a hardware monitor, a session monitor, and a status monitor, as well as control and help facilities, and it can also manage equipment and applications.

However, NetView still doesn't fit the definition of the ideal net management system. The reason? It mainly integrates the management of a Systems Network Architecture network with IBM equipment residing on it. When other vendors' equipment exists on the network, NetView becomes an alarm collector and reporter, and even then only through NetView/PC and only with those vendors that conform to IBM's specifications.

IMPLEMENTATION AND DESIGN

The corporate network's ability to connect various end users (user to user, application to application, or user to application) and to carry information between them consistently is critical to a corporation. Therefore, successful management of the corporate network is as important as any other corporate function. In order to maintain the network's service levels, the user must manage it as efficiently as possible.

Network management can be implemented using the tools and facilities available today and still be functional until standards arrive. To implement a management system, the user must first manage the physical part (end-to-end physical connectivity) and then focus on the logical part (end-to-end logical connectivity). If these two tasks are completed successfully, the user will be able to efficiently manage the current corporate network, and implementing the forthcoming standards will not require redesigning it.

The first step in designing a complete net management system is to implement a physical network management system. This involves, among other tasks, an inventory of all network components down to the end-node level. An end node is any device on the network that has more than one input or output; for example, an IBM or DEC host, terminal server, packet assembler/disassembler, 327X terminal controller, front-end processor, or T1 multiplexer.

During this phase, users should construct a complete, detailed map of the network, showing all physical connections between end nodes, indicating the kind of traffic and protocols used and the available and utilized bandwidth for each link.

When these steps are completed, the user can select and implement a physical network management system. This system will follow the first design approach discussed above, which focuses on network availability control. It will be the basis for the next step toward complete network management.

Examples of a physical network management system are the recently introduced Meridian Data Network System from Northern Telecom, Inc., AT&T's Unified Network Management Architecture (UNMA), and Codex's 9800 series of its Integrated Network Management System.

The second step in the design and implementation of the complete net management system is the construction of a logical system. Once the physical network management system is in place, logical system construction will become a matter of identifying the applications that reside on each host, the access methods they support, and the paths they utilize.

To construct a logical net management system, the user must inventory all the applications accessed through the network, including descriptions of the connection methods and protocols supported, such as dial-up connections, dedicated access, asynchronous, Synchronous Data Link Control, or X.25.

The user must also perform a user population study to determine user locations, types and methods of connections, and the usage of the network in general.

Homegrown Network Management

When these steps are completed and the data collected is added to the network map, a logical network management system can be selected and implemented. This system must complement and be compatible with the physical net management system. Examples of currently available logical management systems are NetView and Cincom Systems, Inc.'s NetMaster.

At this point, assuming that the system designer followed the guidelines specified previously, the integration of the physical management system and logical management system can be accomplished.

For example, if the systems selected are compatible, reports on path availability or path failures from the physical system can be directed to the logical system, which can then adjust pacing so the active sessions do not overload the remaining physical paths.

If the logical system detects slow response time or a high number of retransmissions during a session, it can inform the physical system, which in turn can reroute the session over a new path. Since the logical system is aware of the status of applications, it can, for example, notify the physical system when sessions should be restricted to a specific physical path (such as an encrypted link), thus preventing any security breaches.

From this point forward, the effectiveness of the management system will depend mainly on how well the network operators are trained and the accuracy of the data maintained for all applications and users. This will make any transitions or conversions to a future Open Systems Interconnection-based system almost painless.

Even if the user chooses not to convert, the benefits from efficient corporate network management will be numerous.

FUTURE DIRECTIONS

A true net management system provides a common interface between all the various architecture- and vendor-specific NEMSS. Its functionality and structure should be flexible enough to allow it to be interfaced with any vendor's NEMS with minimal effort. The ideal net management system will provide a common user interface for the network operator, can run under any operating system or hardware platform, and can be interfaced to and manage any kind of network. Its functionality should at least include the facilities mentioned in the International Standards Organization (ISO) OSI model. Unfortunately, the ideal doesn't exist commercially.

In recent years, many attempts have been made to provide the ideal net management system. All the vendors with a networking strategy have announced a management system, based on their environment, and have claimed it to be the de facto industry standard.

Current limitations and vendor insistence on proprietary architectures notwithstanding, the marriage of the two schools of thought outlined above is imminent. Based on the progress of the standards bodies under the ISO/OSI architectural umbrella, true net management systems will be arriving in the near future.

In the ISO/OSI model, the network management aspects are covered under the Draft International Standard 7498/4 (addendum to the basic ISO model definition), the Draft Proposal 9595 (management information services) and Draft Proposal 9596 (management information protocols). While the details and protocols defined in these drafts are beyond the scope of this report, they can be summed up in a few words.

The ISO network management specifications define several protocols that encompass the basic layers of the OSI model. These network management protocols do not constitute an eighth layer but rather provide adjunct functions to the other layers for data collection and information distribution. The data collected and distributed deals with fault, configuration, accounting, performance, and security management, as well as naming and addressing.

The area of network management standards is the most difficult to implement and, unfortunately, has not received the proper amount of attention until recently. Due to the realization that a networking architecture is not complete without net management system specifications, the ISO's Joint Technical Committee 1 SC21, or OSI committee, has amplified its efforts, and more people worldwide are engaged in standards development.

AT&T's UNMA also addresses the problem. It even claims to use the ISO/OSI model as its architecture, but the fact remains that it is a system provided by a vendor and other vendors might not choose to adhere to it.

When the ISO/OSI model becomes an industrywide standard supported by the user community, vendors will have to adhere to it. The management of the corporate network will be accomplished from a single location, or if needed, from distributed locations, using a common operator interface and an international set of standards. □

Catching Up to the Future of Integrated Network Management

This report will help you to:

- Evaluate the probable changes in network management systems that should occur within the next three years.
 - Recognize the driving forces behind evolving network management systems.
 - Understand the benefits of end-to-end network management via a single system.
-
-

Assessing the future of network management systems is more than an intellectual exercise. Network managers must make decisions that have substantial dollar and resource costs attached to them. These decisions should be based not only on today's knowledge, but on future resource availability as well.

As we discuss possible solutions to the limitations of today's net management systems, we're ignoring concepts and technologies that are only vaguely understood today.

It's unlikely that some totally new concept or technology will emerge over the next three to five years that will have a substantial impact on telecommunications management, but network management will undergo some definite changes in the near future.

FULL NETWORK CONTROL

The most important change in network management will come from the development (albeit slowly) by 1992 of net management systems that can truly manage an entire network through a single system consisting of products from a number of vendors.

This Datapro report is based on "Catching Up to the Future of Integrated Net Management," by Dennis Krentz, Market Analysis Co., from *Network World*, February 1989. © 1989, NW Publishing, Inc. Reprinted by permission.

The following factors will drive the progress toward this integrated system:

Market Demand

In the early stages of deregulation, network managers and other telecommunications decision-makers were happy simply to have a growing range of equipment and control tools from which to choose. Demand for integration began to increase substantially by the mid-1980s, and many surveys show that the primary concern of today's telecommunications managers is the net management system's inability to deal effectively with all of a typical network's components. By mid-1989, individual vendors will begin announcing development efforts that emphasize total system management capabilities.

Development of Industry Standards

The current move toward development of standards, in the form of either products (such as IBM's NetView and Timeplex's TimeView) or protocols (such as Open Systems Interconnection), will accelerate as clients put more pressure on vendors to speed up the process.

Catching Up to the Future of Integrated Network Management

Alliance Development

As the focus on industrywide standards and market pressures from the buying community increases, individual net management system vendors will attempt to achieve short-term integration objectives through agreements with indirect competitors. These agreements will initially involve vendors that do not directly compete with one another (such as subrate modem, private branch exchange, and T1 vendors).

However, by late 1989 or early 1990, market demand pressures will cause a few major vendors to announce exchange-of-information programs in anticipation of future agreement on one of the industry standards currently under development.

Enhanced Interface Development

Although a number of tools currently exist for constructing interface capabilities between divergent control systems, they are hampered by two general problems: vendors have no strong need for them, and they are unwilling to share the information necessary to use the tools effectively.

The combination of increasing market demand for integrated systems, industry acceptance of universal standards, and the initial movement toward information sharing will increase the use of interface tools as a stopgap integration measure. This, in turn, will lead to the production and marketing of a series of equipment and system interfaces by mid-1990. Industry-wide standards development will then further the integration process.

Progress toward true integration of monitoring and control capabilities has already begun and will be essentially complete by the mid-1990s. By selecting products carefully, network managers should be able to build a fully integrated voice/data system involving a number of different vendors' products by as early as 1994.

CONFIGURATION CAPABILITY

Increased attention to the management requirements of multivendor networks containing diverse transmission types will necessarily focus on configuration issues within those networks. The current lack of end-to-end network management capability makes the need for improved configuration capabilities somewhat moot.

The ability to monitor and control an entire network through a single system, however, will quickly bring up the need to automate network inventory and de-

sign capabilities. Such capabilities will develop rather quickly and evolve from two currently available tools: expert systems technology and scanner technology.

Attempts to use expert systems technology for network design and configuration have already been made. These early efforts, however, have been largely rules-based rather than inference-based. In rules-based systems, decisions are made based on the rules established within the system; therefore, decision quality reflects the quality of the rules involved. Inference-based systems take this process one step further. The system draws conclusions from a series of data and generalizes from sets of statistical samples.

The possibilities of an inference-based system are especially interesting in light of the rapid advances in data base management technology. Today's data base management systems can supply substantial amounts of computing power at the workstation or personal computer level. This permits users to process very large data sets containing a multitude of details concerning individual transmissions, over time, within a given network.

Today, that information is collected only within subsections of individual networks—a shortcoming directly related to the lack of end-to-end monitoring capabilities in today's network management systems. Net management systems that manage only a portion of a network can collect data only from that portion.

A systemwide data base, combined with an expert system containing an inference engine, provides a path to automated network design capabilities. The network manager can construct a system that will set initial design parameters based on a combination of cost, transmission type, and traffic projections. The network can then be reconfigured almost at will based on the criteria of that particular network (such as volume, transmission speed, least-cost routing considerations, and call priorities), subject to actual physical constraints.

A few network management system vendors are currently considering similar design and reconfiguration systems. The availability of end-to-end monitoring systems will accelerate the production of design and reconfiguration systems, most likely by 1991.

EXPERT DIAGNOSTICS

Expert systems technology has applications throughout the network management process. While it is unlikely that expert systems will have total control of

Catching Up to the Future of Integrated Network Management

large networks during the next 50 years, diagnostics of transmission-related problems are likely to be handled increasingly through automated methods rather than through the largely manual methods used in most networks today.

SAVING TIME AND MONEY

In today's environment, the operator or user must identify the fault within the system, then trace the fault and repair or replace the malfunctioning piece of equipment. This process can take days to complete, resulting in system downtime costs that can run to hundreds of thousands of dollars per hour.

Inference-based expert systems offer the potential for both time and dollar savings in the diagnosis of system faults. In such a system, action can be taken on the problem almost immediately, in contrast to systems that require the operator or user to first recognize an alarm condition (or other method of notification).

The diagnostic procedure within the expert system can be based on a comprehensive routine, both rules-based and inference-based, rather than being limited by the capabilities of either individual pieces of equipment or individual people.

With expert systems, the diagnostic process begins and ends more rapidly, resulting in both time savings and direct cost savings because of decreased system downtime and decreased resources needed for diagnostic activity.

The primary impediment to expert systems today is the inadequacy of data base management resources to process the information required by the inference engine. However, the two technologies are developing rapidly, and at least one vendor is expected to announce the inclusion of such a system in its network management system sometime this year.

REAL-TIME INVENTORY CONTROL

Inventory control within a typical corporate telecommunications network presents the manager with a series of monitoring and design problems. A network cannot be monitored or configured properly if the network inventory is inaccurate. Reliance on manual reporting methods has frequently proven to be ineffective.

Some vendors (notably IBM/Rolm Systems Division) have implemented controls whereby alarms are generated when a piece of equipment is removed from

the system. Other vendors are able to map major components and trunks within the telecommunications system. However, at this time, vendors cannot track system changes completely and automatically.

While a number of manufacturing applications use scanner technology for inventory control, a direct transfer of principles from the manufacturing environment to the telecommunications environment is not likely to occur. Communications devices are typically so geographically scattered that scanning is not feasible. However, aspects of scanner technology and major portions of the process involved are being used to develop comprehensive automated inventory control systems for telecommunications networks.

Inventory update frequency is the primary difference between the two methods of automated inventory identification likely to be in place in some systems by early 1990.

The first method, periodic polling, is similar to that used in automated shop inventory control. Devices on the system are queried periodically, either through a standalone process within the network management system or as part of a larger overhead control process in the system. The devices are identified through a coding process similar to that used in bar coding, except the identification code is located internally to the unit involved rather than on an external label. The host processes the information, which is ultimately used in the network inventory management system.

From periodic polling, two possible methodologies will emerge. The first will be a stand-alone process, seen as a window in the network management system, which will allow devices to be polled periodically. Since polling must take place over the network, polling capability is therefore contingent upon available bandwidth. The second method will include polling as part of some larger system overhead process that will allow devices to be polled as part of a larger control process.

The second method of automated inventory identification is continuous inventory updating. In this process, the device sends continuous or almost-continuous information regarding information type and location to the host as part of the regular system overhead. Because the process is part of system overhead, it must necessarily be a part of the network management system. A continuous (as opposed to periodic) update process may be necessary in systems that are relatively dynamic. As in the periodic polling process, the host receives standardized information, which is then factored into the network inventory management system.

Catching Up to the Future of Integrated Network Management

BANDWIDTH PARTITIONING

In its most basic form, virtual networking is simply the partitioning of available bandwidth with simultaneous implementation of security measures across that partitioned bandwidth. To the user, the partitioned bandwidth appears as a separate circuit (as in a virtual WATS line) or a collection of circuits (as in a virtual network).

In private networks, the challenge is to bring virtual networking capabilities down to low levels while still permitting access and security controls to be maintained over what would now be two (or more) separate networks. The challenge does not rest in the partitioning process itself but in controlling access to partitioned bandwidth. This can be met only through a combination of technical advancements and security arrangements.

Initially, virtual networking will take place through restricted access to specific ports at individual nodes. That is, an individual user will, based on access entitlement, be limited to a series of specific paths through specific ports at individual nodes.

After the virtual network is defined at the port level, access to that network will be restricted according to a rules-based system that relates user access privileges to port mappings. While integrity between different virtual networks and different users is maintained, some problems exist with this method. Most important, unused bandwidth in one virtual network cannot be used by those with access privileges to another network. The result is that at least some of the dynamic routing/rerouting capabilities in today's net management systems will be lost.

The second generation of small-scale virtual networking products, which allocate bandwidth to individual virtual networks through a software-controlled process, will solve that problem. This process will combine dynamic routing capabilities with expert systems technology to allow individual virtual networks to expand or contract given specific geographical and bandwidth requirements. This capability may be available in two or more net management systems by late 1989.

Because the expansion and contraction of individual virtual networks can be handled through a rules-based expert system, and because virtual networking has been available on a larger scale within the public network for some time, little in the way of new technology must be developed. Rather, the technology must be downsized and adjusted before the product can be brought to market.

PARTITIONING POSSIBILITIES

Implementation of bandwidth partitioning capabilities will have a dramatic impact on all telecommunications users, including the home computer and small network management system user.

The primary impact of bandwidth partitioning will be to free up certain resources and provide overall system flexibility. For users and managers of large systems, the first result will be a consolidation of physical networks. Where multiple physical telecommunications facilities are currently required, due to either application constraints or security needs, the creation of a small virtual network within the main network would cancel the need for other physical networks.

Separate but infrequently used networks such as emergency services networks or networks transmitting classified data could be constructed within the larger primary network.

Users with large network management systems currently have other options to attain that flexibility, albeit at a higher cost. Small-system and home computer users, however, generally have no other alternatives simply because they are dealing with lines rather than with a network. Bandwidth partitioning (or the creation of small-scale virtual private networks) can effectively change one or more telephone lines into a small network.

As one example, Pacific Bell experimented recently with an offering (dubbed Project Victoria) that would enable home computer users, through a combination of bandwidth partitioning and multiplexer technology, to receive a combination of informational services, cable television and enhanced telephone services over the same wires through which only basic telephone service is currently available. Similar benefits would accrue to small-business clients within the public network.

The implementation of bandwidth partitioning capabilities also promises to have substantial impact on small net management systems. The ability to partition bandwidth and effectively create smaller subnetworks within small networks will greatly increase the number and type of applications possible within a given network, however small. Applications that either cannot run or must run in series due to bandwidth constraints will be available to system users as bandwidth partitioning becomes more readily available.

Catching Up to the Future of Integrated Network Management

PUBLIC/PRIVATE INTERACTION

As recently as a few years ago, network management systems were structured almost entirely as client premises products rather than network-compatible products. With client premises products, the only interaction with the public network was in using individual trunks provided by public network vendors. With network-compatible products, interaction occurs between the control features within both the client premises and public network equipment.

Digital access and cross-connect system compatibility (where client premises T1 systems can drop and pick up individual T1 channels within the public network) blurred this distinction somewhat. The ongoing development of Integrated Services Digital Network capabilities and the resulting ISDN-compatible claims have blurred the distinction even further. Recent petitions by some Bell operating companies to remove the information-provision restrictions imposed by divestiture indicate they are willing to provide further interaction with private networking systems.

The most feasible (and potentially the most profitable) combination of public and private network control technology in the next two years is the extension of network compatibility to the point where the private network can exert a form of software control over portions of the public network. This is not a new concept—some carriers have been offering forms of software-defined networks (SDN) for at least 18 months.

Further, the most basic premise of ISDN is the ability of the ISDN user or manager to control parts of the public network in much the same way a private network is controlled. From a viewpoint of cost and system management, however, a combination of public network SDN and private network management capabilities makes the most sense in the short term. The private network sector is demanding products that will expand the ability of the network manager to view and configure the network.

The ability to exert software and data base control over portions of the public network at the switch level will allow the private network manager to create or delete additional network capacity as bandwidth demand warrants by adding or deleting circuits within the public network. Effectively, the network manager will be able to create virtual subnetworks or temporary additions to the primary network, which will have a considerable impact on network management practices.

Geographical areas that currently do not warrant fixed network facilities will be integrated into the

master private network by adding virtual subnetworks derived from the public network.

Finally, the availability of expanded networks to individual companies will mean a combined increase in both informational flow and application availability.

Developments and changes in network management systems are becoming increasingly market-driven. As private network managers become able to define the new features that their systems require to function properly, these features will be developed by individual vendors. That development process is expected to proceed more rapidly over the next 18 months than was previously anticipated.

MULTIVENDOR INTERFACE OPTIONS

The lack of integration between different vendors' software and hardware is the primary barrier to the development of a true network management system. Control code and interface information within individual net management systems, or between components within a single vendor's network, is typically proprietary to that vendor.

Individual systems simply do not communicate well with one another. In fact, they cannot communicate unless vendors are willing to either exchange proprietary control, interface and alarm information, or to agree on a single set of standards for that information.

Although a growing number of systems and components can interface through third-party monitoring systems such as IBM's NetView, such interfacing is typically limited to passing status or alarm messages to the host.

The result is that no network management system today can truly manage the components of multiple vendors. This is the result of vendors having no incentive to exchange proprietary information with their competitors. There are, in fact, a wide range of options for constructing some form of interface between individual systems and components.

Current Interface Options

The application program interface (API) takes a particular application within the network management process, performs a software protocol conversion on it and hands it off to another set of controls. The main advantage to this method is that the user can tailor the API to fit a particular situation. A disadvantage is that large numbers of applications require large numbers of APIs.

Catching Up to the Future of Integrated Network Management

The asynchronous ASCII terminal interface is another option. This is the most commonly used of all interfaces. The vendor simply enables the user to view or control the system or component through an ASCII terminal. The extent of the user's ability to view or control is typically limited by the vendor, and enhancements to existing capabilities may be difficult or impossible.

A third interface option is reverse engineering. Control codes and other organizational factors are broken down into individual components and rewritten in a language compatible with the primary system. This method is efficient in that all aspects of the integration process can be addressed. Reverse engineering can be tremendously expensive, though, and even minor system changes that require reverse engineering can mean that the entire process must be repeated.

Software bridging is another interface option. A software bridge passes an application from one system to another, with the passage being transparent to the user. This tool is very efficient from an engineering standpoint, but designing a bridge requires substantial proprietary information from each system.

Standard message protocols are a fifth option. Examples include CCITT X.400 and the International Standards Organization's (ISO) Common Management Information Protocol.

The main advantage of these protocols is that they allow users and vendors to focus on a known set of standards. However, these standards are still under development and are unlikely to be completed until the mid-1900s.

Another option is third-party "umbrella" systems, in which a single set of networking standards is used across a series of networks and equipment. This method has the most promise for the future but requires agreement from a large number of vendors to be successful.

Wrap boxes are the final interface option. These are devices that either bypass or bridge individual components within a system so that the performance of those components can be monitored to some extent. However, rarely is either complete monitoring or control possible. □

OSI-Based Network Management

This report will help you to:

- Compare IBM's NetView with OSI-based network management systems.
 - Examine basic concepts and terms of OSI management.
 - Make effective OSI management decisions.
-
-

During the last three years, network management has become a vital element in voice, data, and image communications. Once considered an option, it is now a necessity for administering and operating both public and private networks. In response to this need, the International Organization for Standardization (ISO) is defining a set of Open Systems Interconnection (OSI) standards for network management.

OSI Management standards are documents written by OSI committees that define information structures, services, and protocols required for OSI Management. These standards define the software tools necessary to monitor, control, operate, and administer network components in OSI environments.

The services provided by these standards are arranged into two broad groups: Common Management Information Services (CMIS) and Specific Management Information Services (SMIS). CMIS standards provide software tools called "primitives," which are the basic means of supporting OSI Management functions. Typical primitives would include get, set, event report, and action. SMIS standards provide additional software tools called "directives." These directives provide specific OSI Management functions

This report was prepared for Datapro by John J. McCann, president of Comnet Systems, a Ridgewood, New Jersey consulting firm specializing in network management. Formerly a Senior Analyst with the international consulting firm Arthur D. Little, Mr. McCann now provides network management implementation studies for *Fortune* 1000 clients. Additionally, Mr. McCann does network product development for carriers and vendors.

within the areas of configuration management, fault processing, performance monitoring, security, and accounting. Additional OSI Management standards such as the Management Framework, the Structure of Management Information (SMI), and the Directory complete the OSI Management standards.

OSI Management standards are important because they offer the only realistic and workable method for enabling networks to exchange management information on a worldwide basis. As public and private networks became more closely integrated in future years, the exchange of administrative and operational data will become increasingly crucial. OSI Management standards will provide each network and network component with an identical set of tools to support information exchange. Without these standards, every network would require proprietary gateways to every other network in order to exchange management information. This is not acceptable on either a cost or an efficiency basis.

Index to This Report	Page
OSI Management Framework	104
OSI Systems Management	105
Specific Management Functional Areas ...	110
Structure of Management Information (SMI)	115

OSI-Based Network Management

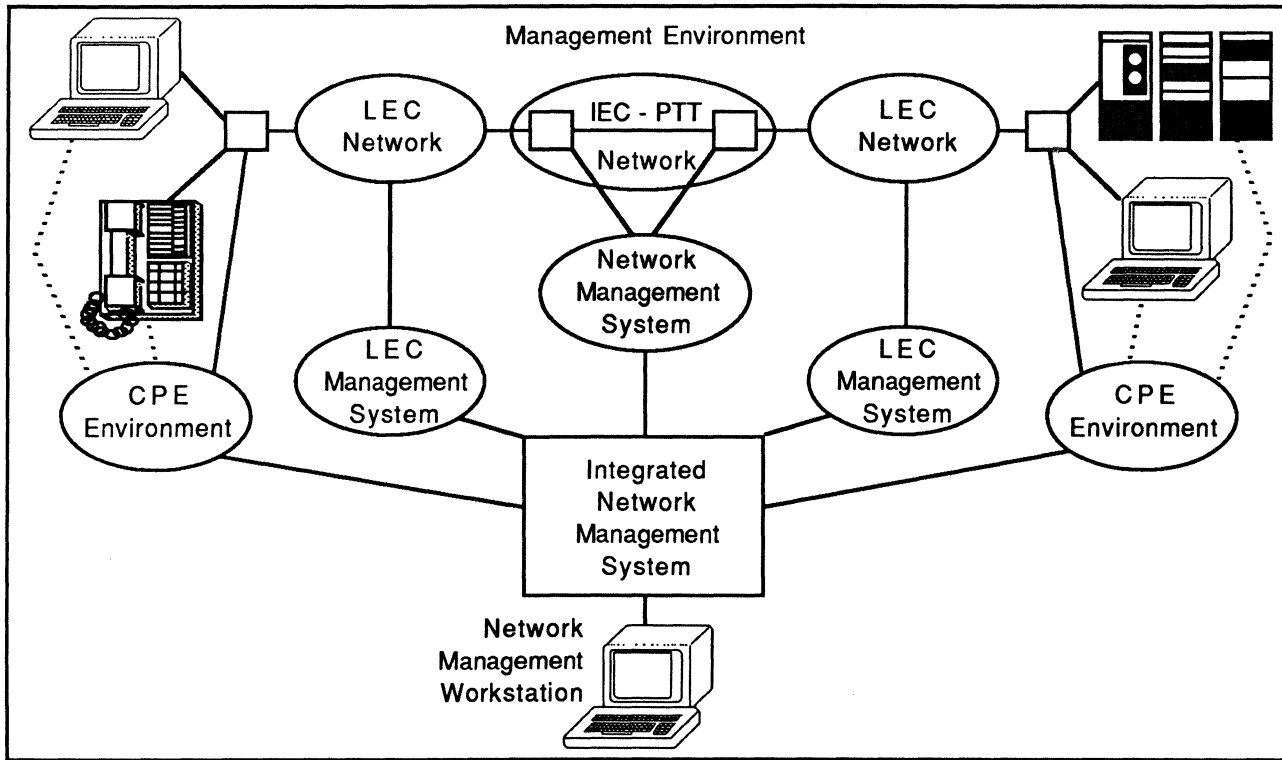


Figure 1. Network management is distributed across three environments: customer-premise equipment (CPE), local exchange carrier (LEC), and interexchange carriers (IECs) and foreign postal, telegraph, and telephone agencies (PTT). Effective end-to-end management requires close cooperation between these three environments.

OSI Management standards will *support*, rather than replace, proprietary management architectures such as IBM's SNA/NetView, AT&T's UNMA, Digital Equipment Corporation's Enterprise Management Architecture (EMA), and Hewlett-Packard's OpenView. For at least the next 10 years, these proprietary management architectures will continue to operate and administer their respective environments while utilizing OSI standards only to define, format, and exchange management information. This will both protect and enhance the value of installed management systems.

Figures 1 and 2 illustrate the scope of network management and further emphasize the importance of OSI Management standards. As shown in Figure 1, network management is distributed across three environments: customer-premise equipment (CPE), local exchange carrier (LEC), and interexchange carriers (IECs/PTT). Effective end-to-end management therefore requires close cooperation between these three environments with their diverse authorities and technologies.

Figure 2 focuses the problem into greater detail. This "composite network" enlarges the CPE environment, showing typical components in *Fortune* 1000 voice, data, and image networks. Figure 2 also shows inter-

faces to the LEC and IEC/PTT switching centers as well as interfaces to Common Channel Signaling System 7 (CCSS7).

How Does OSI Management Solve These Problems?

OSI Management standards are important because they provide uniform software tools to efficiently operate and administer voice, data, and image networks in complex CPE, LEC, and IEC/PTT environments.

Figure 3 shows how OSI Management standards can be implemented for products which reside in LEC or IEC, as well as CPE environments. As illustrated, proprietary products (such as NetView or AT&T's UNMA product, the Accumaster Integrator) located in LEC or IEC switching centers could use OSI Management software tools to exchange operational or administrative data with similar products located on the customer premise equipment (CPE). The interface used in Figure 3 is a Basic or Primary rate ISDN D-Channel. Tariffed network management services, conceptually similar to this diagram, will be offered by carriers in 1989.

OSI-Based Network Management

When will OSI Management products be available? In order to answer this question reasonably, it must be noted that the OSI Management standards discussed in this report are merely OSI documents—they are not OSI products. Significant worldwide efforts are required to develop products based on these standards. Although several vendors have already developed products closely oriented to these standards, true OSI Management products that conform with final international standards will not be available until between 1990 and 1992. Figure 4 provides a schedule for completion of OSI Management standards including Draft Proposals, Draft International Standards, and International Standards.

OSI Management standards, although not yet finalized, can still play a central role in an organization's long-term network management strategy. Since OSI manages "objects" with multiple "attributes," a typical *Fortune 500* corporation can use these broad generic categories to plan long-term management for all its networks. These networks may transmit voice,

data, and image, using both public and private network facilities, with proprietary architectures such as SNA or Digital's DNA. Various switching techniques such as Digital Dataphone Service (DDS), Private Branch Exchanges (PBXs), Local Area Networks (LANs), T1 multiplexers, and packet systems are employed.

By using the specific OSI services provided by the configuration, fault, performance, security, and accounting Specific Management Functional Areas (SMFAs), an organization can plan a unified approach to managing its multiple diverse networks. Critical to this approach are the unified syntax, database, naming, and addressing functions provided by OSI's SMI and Directory standards.

An organization must, of course, develop its network management strategy within the limitations of real-world network environments. User service levels, staffing problems, and cost containment must be paramount issues. It is possible to factor these real-world

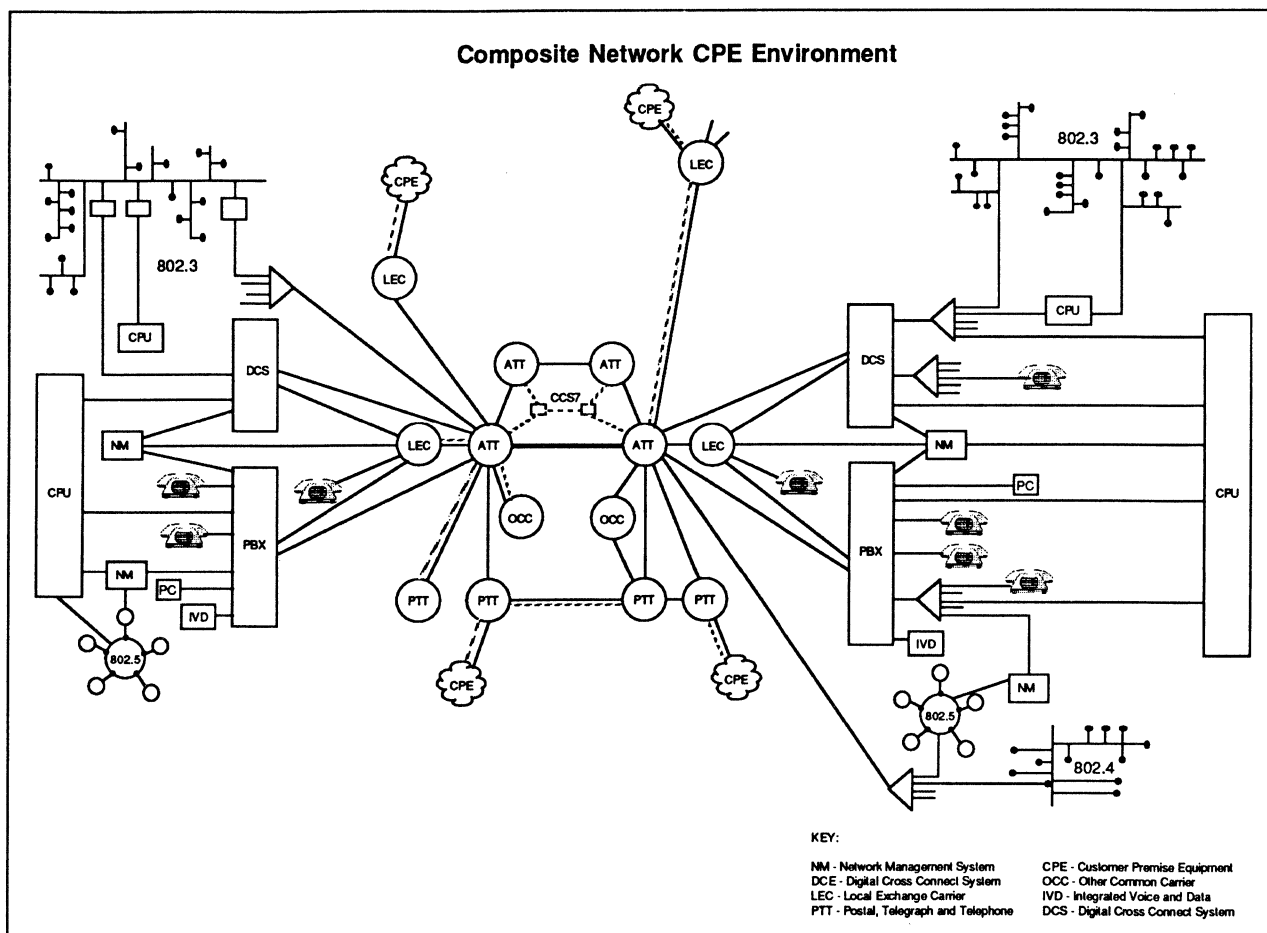


Figure 2. This "composite network" enlarges the CPE environment to show typical components in voice, data, and image networks. This illustration also shows interfaces to local exchange carrier (LEC) and interexchange carrier (IEC)/postal, telegraph, and telephone agency (PTT) switching centers as well as interfaces to Common Channel Signaling System 7 (CCSS7).

OSI-Based Network Management

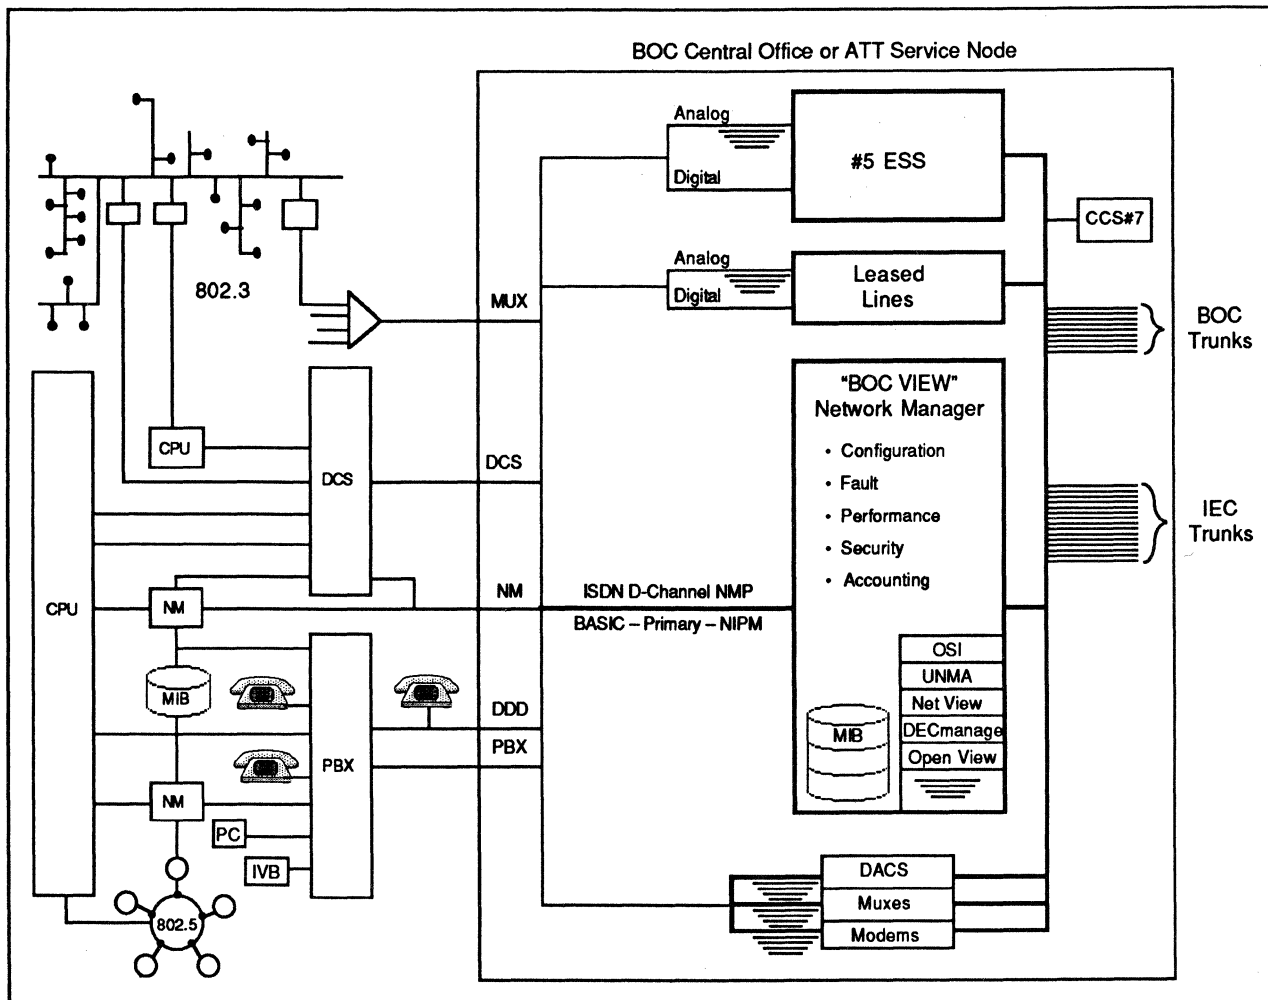


Figure 3. This is one possible implementation of OSI Management standards for products residing in local exchange carrier (LEC), interexchange carrier (IEC), or customer-premise equipment (CPE) environments. Proprietary products located in LEC or IEC switching centers, such as NetView or AT&T's UNMA product called Accumaster Integrator, could use OSI Management software tools to exchange information with similar products located on the customer's premises (CPE).

requirements and OSI Management tools into a cohesive management strategy. To do so, a number of discrete steps are required. These steps integrate related issues such as installed networks, planned networks, corporate business objectives, cost/chargeback policies, available management products, and Network Management Center (NMC) staffing requirements.

With a proper approach and reasonable objectives, this strategic analysis should provide unified network management, along with attendant benefits of increased productivity, wider dissemination of information, and better control of service and procurement costs.

OSI MANAGEMENT FRAMEWORK

The OSI Management Framework, DIS 7498-4, is an ISO standards document which establishes guidelines for coordinating the development of existing OSI Management standards.

The OSI Management Framework serves as a reference document for other OSI Management standards. The Framework:

- defines the terminology of, and describes concepts for, OSI Management;
- provides an abstract model of OSI Management and gives an overview of OSI Management's objectives and facilities; and
- describes OSI Management activities.

OSI-Based Network Management

The Framework functions in an OSI Management environment—a subset of the total OSI environment. This subset encompasses all tools and services needed to control and supervise interconnection activities and managed objects. The OSI Management environment includes both the capability for managers to gather data and exercise control and the capability to maintain an awareness of, and report on, the status of managed objects.

The major issue defined by the Framework standard is OSI Management *facilities*. The term *facility* is not formally defined within the Framework standard; however, it can best be described as a set of functions that accomplish specific objectives. The functions are grouped into the following five categories.

- Configuration
- Fault
- Performance
- Security
- Accounting

The Framework also defines the structure of OSI Management within the following three groups:

- **Systems Management** provides mechanisms for monitoring, controlling, and coordinating all managed objects within open systems;
- **Layer Management** provides mechanisms for monitoring, controlling, and coordinating each of the seven layers in the OSI Reference Model (for more information, see Report CMS20-010-201, "The ISO Model for Open Systems Interconnections"); and
- **Protocol Management** provides mechanisms for monitoring and controlling a single communications transaction.

The Framework standard introduces the concept of the Management Information Base (MIB). The MIB is "that information within an open system which may be transferred or affected through the use of OSI Management protocols." The MIB has all information related to managed objects within the OSI environment, whether these objects are software modules, PBX switches, analog lines, T1 multiplexers, state variables, telephone sets, personal computers, or any of a thousand entities used in data communications systems.

It is important to note that the MIB concept does not imply any form of physical or logical storage of infor-

HOW THE MANAGED OBJECTS ARE MANAGED

The five Specific Management Functional Areas (SMFAs) monitor and control a managed object through four aspects:

- **The Object's Existence**
- **The Object's Attributes**
- **The Object's States**
- **The Object's Relationships**

Existence—A managed object *exists* if it has an object identifier and an associated set of management information that is accessible through OSI Management services.

Managed objects can be created or deleted. To create a managed object, the user places into the MIB the object's identifier and a set of information appropriate to the object's class.

Attributes—describe properties of the object, such as operational characteristics. An attribute has an ID and a value. During the object's existence, only the values can be changed—the attributes themselves cannot be created or deleted.

State—represents the instantaneous condition of the object's availability and operability. For example, a multiplexer's state may be represented as 11, meaning available and operable. Conversely, state 10 may indicate available, but inoperable.

Relationships—define the interdependence between the managed object in question and other managed objects. For example, a relationship exists between an OSI terminal and the OSI packet switch which provides protocol processing and routing for that terminal.

mation. Since the issue of storage is local to the open system, it is considered to be outside the proper domain of OSI.

OSI SYSTEMS MANAGEMENT

OSI Systems Management provides mechanisms for monitoring, controlling, and coordinating all managed objects with open systems. First presented in the Framework DIS, 7498-4, this concept was greatly expanded at the March 1988 SC21 Working Group 4 meeting in Washington, DC.

OSI-Based Network Management

OSI MANAGEMENT STANDARDS		EXPECTED REGISTRATION DATES
Title	Current Status	Int'l. Standard
OSI Management Framework (ISO 7498/4)	International Standard	October 88
OSI Management Information Service Overview (DP 2683)	Draft Proposal	July 90
Structure of Management Information (DP 2684)	Draft Proposal	July 90
Common Management Information Service (CMIS) (ISO 9595)	Draft International Standard	September 89
Common Management Information Protocol (CMIS) (ISO 9596)	Draft International Standard	September 89
Configuration Management (DP 2686)	Draft Proposal	July 90
Fault Management (DP 2687)	Draft Proposal	July 90
Security Management (N 2698)	Working Document	July 90
Accounting Management (N 2689)	Working Document	April 91
Performance Management (N 2673)	Working Document	April 91

Figure 4. Current status and expected final registration dates for OSI Management standards.

Working Draft (WD) N2683 defines OSI Systems Management using two major concepts:

- System Management Models
- Systems Management Standards

Additionally, WD N2683 introduces general conformance requirements for OSI Management standards. Conformance requirements specify the elements which must exist in any offering claiming to be an OSI Management product.

Systems Management Models

OSI Management models define various aspects of Systems Management. (See Figure 5.) Also, these models provide a conceptual and terminological framework for:

- Common Management Information Services Element (CMISE);
- Specific Management Functional Areas (SMFAs);
- Structure of Management Information (SMI); and

- Generic Definition of Management Information (GDMI).

Working Draft N2683 presents three conceptual models to define Systems Management. These are:

- Functional Model
- Organizational Model
- Information Model

The Functional Model introduces the concept of Specific Management Functional Areas (SMFAs). Prior to the March 1988 meeting in Washington DC, these SMFAs were referred to as Specific Management Information Services and Protocols (SMIS/SMIP). ISO has defined five SMFAs:

- Configuration;
- Fault;
- Performance;
- Security; and

OSI-Based Network Management

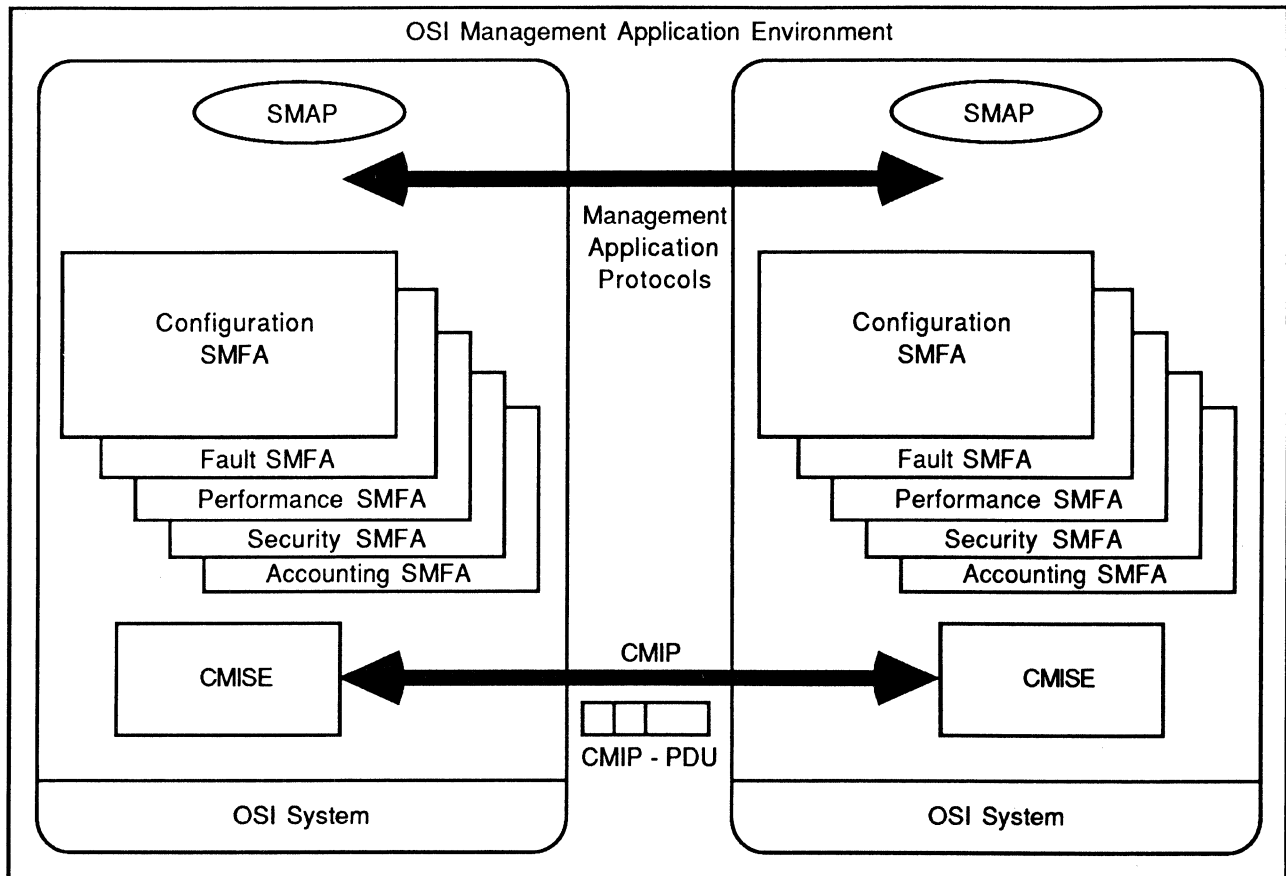


Figure 5. This figure illustrates the components of the Application Layer involved in the interconnection between peer management application processes.

- **Accounting.**

Each SMFA is defined in a separate OSI Management standard that defines the set of facilities to support the SMFA functions; procedures associated with these facilities; use of CMIS to provide these facilities; classes of managed objects within SMFA; and subsets of facilities to provide conformance classes.

The organizational model describes the distributed nature of OSI Management through the use of the following concepts:

- Managed Open System;
- Managing and Agent Processes; and
- Domains.

This is an important model since it relates abstract OSI concepts to the real world where management centers, systems, and people are widely distributed. It basically defines how large organizations can distrib-

ute management control across the CPE, LEC, IEC, and PTT domains. Figures 1 and 2 illustrate these concepts.

The Information Model defines the concept of Managed Objects. (See glossary sidebar.) It specifies their attributes, the operations that may be performed upon them, and the notification that they may issue. The set of managed objects in a system, together with their attributes, constitute that system's Management Information Base (MIB).

System Management Standards

The System Management Working Draft provides a brief overview of the remaining OSI management standards which include:

- **CMIS/CMIP**—services and protocols for use in the invoker-performer dialogs that support the OSI Management functions;

OSI-Based Network Management

OSI GLOSSARY

Directory—a collection of open systems that cooperates to hold a logical database of information about a set of objects in the real world.

Directory Information Base (DIB)—the set of information managed by the Directory.

Directory Information Tree (DIT)—representation of directory entries within the DIB. Each DIT entity has a distinguished name, which unambiguously identifies that entry. Entries higher in the tree (nearer the root) may represent countries or organizations, while entries lower in the tree may represent people, application processes, or OSI-defined elements.

Directory Schema—the set of rules that ensure the tree structure is maintained within the directory.

Directory Services—Interrogation of the Directory and/or modification of the Directory.

Local Systems Environment—Those resources which exist in a real open system, but which are outside the scope of the OSI environment.

Managed Object—A data processing or data communications resource that may be managed through the use of an OSI Management protocol. The resource itself need not be an OSI resource.

A managed object may be a physical item of equipment, a software component, some abstract collection of information, or any combination of all three.

Management Information—information, associated with a Managed object, that is required to control and maintain that object.

Management Information Base (MIB)—a conceptual composite of information about all managed objects in an open system.

OSI Environment—those resources that enable information processing systems to communicate openly, that is, to conform to the services and protocols of Open Systems Interconnection (OSI).

OSI Management—the facilities to control, coordinate, and monitor the resources which enable communications in an OSI environment.

- **Configuration SMFA**—a set of facilities that identifies, monitors, and controls OSI-managed objects to support continuous operation of interconnection services;
- **Fault SMFA**—processes faults in an OSI environment. The original OSI fault standard included fault correction. Current standards, however, consider fault correction to be either a configuration SMFA function or to be outside the OSI environment;
- **Performance SMFA**—a set of facilities needed to evaluate the behavior of managed objects and the effectiveness of interconnection activities;
- **Security SMFA**—enables the control and distribution of information to various open systems for use in providing OSI security, for reporting on the security provided, and for reporting on any security-related events that have occurred;
- **Accounting SMFA**—a set of facilities that enables charges to be established for the use of managed objects and cost to be identified for use of those managed objects;
- **Structure of Management Information (SMI)**—defines the logical structure of OSI Management

information. It also establishes principles for naming managed objects and their attributes. SMI defines a number of subobject types and attributes types that are, in principle, applicable to all classes of managed objects; and

- **Generic Definition of Management Information**—there is no consensus within OSI that GDMI should be a separate standard. This document is therefore not being reviewed.

Common Management Services and Protocols

Common Management Services and Protocols support the exchange of information and commands between peer applications in open systems for the purpose of systems management. Common Management Services and Protocols standards and related standards include:

- CMIS-N2835—Common Management Information Services;
- CMIP-N2836—Common Management Information Protocols;
- ACSE-ISO8650—Association Control Service Element; and

OSI-Based Network Management

- ROSE-ISO 9072—Remote Operations Service Element.

ACSE and ROSE standards are not officially part of OSI Management standards. They are, however, essential to OSI Management services and are included here to provide a more complete picture of the OSI Management environment.

Common Management Information Services (CMIS)—CMIS is an Application Service Element (ASE) which is used by an application process to exchange information and commands for the purpose of Systems Management. Basically, this standard defines a set of service primitives that constitute the ASE, as well as the parameters passed in each primitive. CMIS also defines any information necessary for proper description.

The following 10 CMIS service primitives form the basis for virtually *all* OSI Management activities:

- **Confirmed-Event Report**—report a managed object event to peer; reply
- **Event Report**—report a managed object event to peer; no reply
- **Confirmed-Get**—request information retrieval from peer MIB; reply
- **Confirmed-Set**—request information modification in peer MIB; no reply
- **Set**—request information modification in peer MIB; no reply
- **Confirmation-Action**—request peer to perform an action; reply
- **Action**—request peer to perform an action; no reply
- **Linked-Reply**—request peer to provide correlated replies to multiple requests
- **Confirmed-Create**—request peer to create a managed object; reply
- **Confirmed-Delete**—request peer to delete a managed object; reply

The SMFAs (configuration, fault, performance, security, and accounting) use CMIS services primitives. This report explains how the SMFAs use these CMIS services primitives in a later section.

M-CONFIRMED-SET

InvokeID—ID arranged to this invocation

AccessControl—Provide input to access control function

ManagedObjectID—Managed object class and instance

CMISSYNC—Defines how several operations can be done:
Best Effort—"do what you can"
Ordered—"keep operation sequences"
Stop on Error—"if error, do not proceed"
Atomic—"do everything, or nothing"

CMISFilter—test an entry against predefined limits

ManagementInfoList—information identifiers

CurrentTime—time at which response was generated

Errors—area set aside for multiple error codes

Figure 6. The CMIS Draft International Standard (DIS 9595) defines services in terms of service primitives plus parameters. This figure gives an example of the SET primitives. This example illustrates the power and versatility of OSI Management, as this primitive could be applied to any open systems environment, including PBXs, digital cross connects, T1 switches, etc. This example also indicates how complex OSI products must be to achieve compliance with standards.

Services are defined in the CMIS Draft International Standard (DIS) in terms of the service primitives plus parameters. Figure 6 gives an example of the SET primitives.

The power and versatility of OSI Management is apparent in an analysis of this single CMIS primitive. This primitive could be applied to any open systems environment, including PBXs, digital cross connects, T1 switches, LANs, WANs, etc.

For example, Figure 6 also indicates how complex OSI products must be and points out the problems of conformance when multiple OSI systems must work together.

Common Management Information Protocols (CMIP)

CMIP provides a request/response service between peer users in OSI open systems. It also provides event reporting and event monitoring. Specifically, the standard defines an abstract syntax for CMIP protocol data units and procedures for transmitting management information between peer application entities. Additionally, it defines procedures for correctly interpreting protocol control information and conformance requirements that must be met by CMIP products.

OSI-Based Network Management

Essentially, CMIP provides a written procedure for each CMIS primitive. The procedures, along with the protocol data units for all primitives, provide information necessary to eventually develop products for OSI Management.

CMIP procedures for the confirmed-set primitive example in Figure 6 are as follows:

- The function of the M-set and confirmed-set operation is to request the performing service user to replace the current values of the specified management information elements of the specified managed object. The invoking service user supplies the identification of the managed object and the identification and replacement values of the management information elements it desires to change, along with optional access control, synchronization, and filter parameters, to the M-set or confirmed-set service primitive.

Similar additional procedures are defined in a subsequent clause of the CMIP standard.

The protocol data units (PDUs) are the actual information blocks sent from one application entity to its peer application entity. For example, peer applications running in OSI switching multiplexers will swap PDUs to exchange typical information such as network routing parameters. Contents of the PDU are defined in Clause 8 of the standard and are written in a format called Abstract Syntax Definition (ASN.1).

Association Control Service Element (ACSE)

As defined in ISO 8650, ACSE provides those service elements that are used within OSI Management. These include initialize, terminate, and abort.

The initialize primitive enables a user to establish an association with a peer user. This association is required before any of the 10 CMIS primitives defined earlier can be used. The terminate and abort primitives are used to end this association. Initialize parameters and procedures are covered in detail in Clause 6 of CMIP.

Remote Operations Service Element (ROSE)

As defined in ISO 9072, ROSE provides services to all 10 primitives defined above in CMIS. These include invoke, result, error, and reject services. It might also be noted here that ROSE, in turn, depends on certain services from the Presentation Layer of OSI.

SPECIFIC MANAGEMENT FUNCTIONAL AREAS (SMFA)

ISO is currently finalizing the following SMFA standards. (See Figure 5 for projected completion dates.):

- Configuration Management
- Fault Management
- Performance Management
- Security Management
- Accounting Management

These standards describe how CMIS and CMIP are used to achieve the procedures specified for each SMFA. Specifically, each SMFA standard defines a set of facilities that support the SMFA functions, the procedures associated with these facilities, and how CMIS is used to provide the facilities. Each standard also defines classes of managed objects used in the SMFA, as well as facility subsets that may be used for specifying conformance classes.

These SMFA standards have been organized on a topical basis; that is, they are organized in the most useful manner.

At the highest level, SMFAs delineate how to perform the management tasks that accomplish OSI Management goals. Specifically, SMFA functional requirements make use of specific facilities which, in turn, consist of the procedures and information required to accomplish those OSI Management functions which fall within the scope of the particular SMFA. These specific facilities are provided by the exchange of one or more management application protocol data units (MAPDUs).

A MAPDU is a management command, response, notification, or other communication between open systems. It is conveyed between a managing process and agent process. The MAPDU is the unit of exchange in the highest level peer-to-peer management protocol between two systems. The MAPDU employs a CMIS primitive or other application layer service primitive as the means of transfer between systems. More than one specific facility may map onto the same MAPDU. For example, fault management and configuration management may both need to retrieve error counters using identical M-set primitives, where the managed object is a particular X.25 switching processor and the object attribute is the line 10 CRC error counter.

OSI-Based Network Management

No distinction is made in the management protocol to indicate which SMFA causes a MAPDU transmission. That is, MAPDUs provide no information to indicate the SMFA to which they belong. Specific facilities define parts of the MAPDUs, including the semantics of the information carried in the MAPDUs. For example, fault management may specify that severity levels must be present on all fault management error reports.

ISO intended that SMFAs should not specify identical functionality, particularly since SMFAs are divided topically in a somewhat arbitrary way. Instead, one SMFA should refer to the use of facilities specified in another. For example, in order to provide for fault correction, fault management must use configuration management's reconfiguration facilities.

Requirements for management facilities differ greatly, depending on the equipment, technology, and resources being managed. Modems and simple relays are likely to have vastly different management requirements than large mainframe systems. Similarly, a management system for a single-office PBX will generally differ from one that must manage a large wide area distributed system. Despite the differences, however, there are many aspects of management which are common. For example, most management systems require that faults are reported. Many will require the logging of events and gathering statistics.

In order to select *groups* of facilities, *subsets* composed of specific facilities are defined in each of the SMFA standards. Combinations of these subsets (from one or more SMFAs) define ASEs. For example, an ASE may consist of one subset from fault management, one from configuration management, and one from security management.

Configuration Management SMFA

Configuration management is the monitoring and controlling of normal operations in an open system or network. Configuration management SMFA enables network personnel to set up, observe, and change the operational parameters and conditions that control the minute-to-minute interconnection services of an open network. Such parameters and conditions include, but are not limited to:

- existence, names, and relationships of network components;
- addressing information;
- operational characteristics such as line speed, etc.;

- information on whether components are usable;
- conditions for backup operations; and
- routing control.

Configuration management involves defining, collecting, monitoring, controlling, and using *configuration data*. Configuration data includes any information about OSI system resources needed to manage that system. Configuration data represents both static and dynamic information. System administrators use configuration data in a variety of areas such as inventory management, network design, network configuration/reconfiguration, system generation, operator support, and similar functions.

It is important to note that a real, albeit fuzzy, distinction exists between certain configuration management functions and those functions governed by the remaining four SMFAs (fault management, security management, performance management, and accounting). For example, configuration management does not control, or provide information about, system faults—that falls under the jurisdiction of fault management. Fault management's primary purpose is to provide information about faults; yet, it also provides procedures that control certain aspects of systems operations. Fault management does not provide interconnection services to the end user, although it does ultimately contribute to assuring those services are delivered. If some compensatory action (which amounts to a system configuration change) is needed to correct a fault and restore normal interconnection services, then configuration management is responsible.

Configuration management monitors and controls a managed object through four aspects: the object's existence, attributes, states, and relationships (see "How the Managed Objects Are Managed" sidebar). CMIS primitives are the commands used to convey information about these four aspects.

CMIS primitives are grouped into five categories called *facilities*.

- **Object Configuration Facility**—used to create, delete, and rename managed objects. This facility uses

CREATE-OBJECT	DELETE-OBJECT	SHOW-OBJECT
ENROL-OBJECT	DEENROL-OBJECT	REENROL-OBJECT

Figure 7. These six directives are used to implement the object configuration facility, one of five facilities provided by configuration management. The object configuration facility enables a user to create, delete, and rename managed objects.

OSI-Based Network Management

four CMIS services: create, delete, get, and event report. Six directives are used to implement object configuration (see Figure 7.).

- **Attribute Management Facility**—used to examine and set attributes and report on changes in those attributes. This facility uses three CMIS services: get, set, and event report. Three directives are used to implement attribute management.
- **State Management Facility**—used to examine and set the state of existing managed objects and report on state changes. This facility uses three CMIS services: get, set, and event report. Three directives are used to implement state management.
- **Relationship Management Facility**—used to examine and set relationships among existing managed objects and to report on changes within those relationships. This facility uses four CMIS services: create, delete, get, and event reporting. Relationship management is implemented through seven directives.
- **Software Distribution Facility**—manages software distribution. This facility is not yet completely defined. Eventually, it will use both CMIS and File Transfer Access Management (FTAM) to enable the user to perform the following tasks: request that a peer user transmit a configuration (the software constituting a managed object); provide software to other peer users; examine, update, or maintain software components at a peer location; and request cross-network software loading and/or starting.

Fault Management SMFA

OSI fault management is the set of facilities needed to detect and identify abnormal operations in the OSI environment. Faults may cause open systems to fall short of their operational objectives; they may be persistent or transient. Faults manifest themselves as particular events (such as errors) in the open system's operation, and faults are recognized via error event detection. Fault management facilities maintain and examine error logs; accept and act upon error detection; trace faults; and carry out sequences of diagnostic tests.

The fault management SMFA standard defines a set of facilities supporting fault detection and diagnosis, as well as procedures associated with these facilities. The standard also specifies the CMIS services that support these facilities and classes of managed objects handled by these facilities. The standard also specifies a subset of the fault detection/diagnosis facilities that enable realistic conformance classes.

TABLE 1. OSI ELEMENTS DEFINED IN THE STRUCTURE OF MANAGEMENT INFORMATION (SMI) STANDARD

Attribute	Element	Notification
Class	Gauge	Operation
Counter	Gauge-threshold	Report
Counter-threshold	Hidden Element	Subclass
Current State	Log	Superclass
Defined Event	Log Control	Tide-mark

Network personnel can use fault management information to detect abnormal system operation and direct corrective actions. There are three primary fault management activities:

- **Fault Detection**—faults can be detected in three ways. First, users may detect faults during normal operations by routine monitoring procedures or by error report generation. Second, faults become obvious while performing confidence tests. Third, impending faults may be anticipated when threshold values are exceeded.
- **Fault Diagnosis**—faults can be diagnosed by analyzing error and event reports on managed objects or by performing diagnostic tests. Diagnostic tests exercise a managed object in an attempt to reproduce errors.
- **Fault Correction**—fault correction is the ability to repair faulty managed objects. In general, fault correction is achieved through configuration management and/or by operator intervention.

Fault management supports these three activities through a set of six facilities. (Recall that all SMFA standards possess a set of facilities.) The set of facilities are:

- **Spontaneous Error Reporting Facility**—enables a user to send error reports to other users. This facility uses three directives: initiate error reporting, terminate error reporting, and error reporting.
- **Cumulative Error Gathering Facility**—enables a user to periodically request error counter information from another user utilizing CMIS GET services. With this facility, the user can poll the error counters in a peer-user system to discover error conditions. Directives are provided to get error counters, set error counters, and report error counters.
- **Error Threshold Alarm Facility**—provides the user with several threshold functions: send threshold reports, set error threshold values, and request current threshold settings. An error threshold alarm is

OSI-Based Network Management

an event report which notifies a user that a specific counter or gauge has reached a predetermined value. Four directives are available in this facility.

- **Confidence and Diagnostic Testing Facility**—enables one user to direct a peer user to perform a test on a managed object to determine if it is capable of performing its service or of assisting in diagnosing faults.
- **Event Tracing Facility**—enables a user to instruct a peer user to make a local trace or log of specified events.
- **Management Service Control Facility**—provides basic procedures for controlling management aspects of managed systems. In particular, this facility establishes profiles for systems management services and enables the user to initiate, terminate, suspend, and resume those services. (NOTE: SC21-WG4 working notes indicate that this facility may be better placed in systems management information rather than in fault management. The debate is ongoing.)

Performance Management SMFA

Performance management is the set of facilities needed to evaluate the behavior of managed objects and the effectiveness of communications activities. Its management function is the long-term evaluation of OSI systems. This management function requires gathering statistical data in order to analyze operations trends in the communications between open systems.

Performance management encompasses the set of functions, data messages, and parameters to monitor performance; collects and analyzes system statistics; tunes and controls performance based on these statistical analyses; and provides useful reports (both real-time and offline) on network performance.

To be effective, a network management system must monitor the performance of specified events, measures, and resources. This includes facilities to allow network users to select which objects to monitor as well as the length of monitoring time. It also must allow users to specify the frequency of monitoring and to set or change threshold measurement values. It is important to note that procedures and products must efficiently process the collected data for the network management system to be effective, although this is not strictly within the OSI scope.

When there are indications that performance is degraded or otherwise abnormal, users must be able to

tune the system to provide better performance. To adjust for accuracy, the user must be able to change resource allocations, modify resource attributes, and reassign resource attribute values. The issue is significant uncertainty with the SC21 Working Group concerning the proper agent for these tuning functions. These functions may be properly part of configuration management rather than performance management. The issue is currently under investigation.

A network management system must also provide performance reports at the user's request. These reports may represent current, realtime performance, or they may represent historical network performance which can be used as network benchmarks. The reports should relate to an individual open system (a subset of open systems) or to the network as a whole. Additionally, the network management system should accumulate daily performance reports on a weekly, monthly, and yearly basis. These reports should address function issues such as network utilization versus load, ratio of overhead packets to data packets, traffic details and peak-hour rates, and average response times.

As is true with the other SMFAs (configuration, fault, security, and accounting), the performance monitoring standard defines its facilities in terms of CMIS primitives needed to support these facilities. These primitives include event report, get, set, and action. The pertinent CMIP protocols are used.

Additionally, performance management SMFA functions are grouped into two categories: monitoring and control of monitoring.

- **Monitoring Functions** govern synchronous and asynchronous information transfer, periodic information transfer, and notification of a communications instance. Monitoring functions also measure round-trip delay time.
- **Control of Monitoring Functions** allows the user to alter thresholds, events, system operation modes, traffic generation and protocols, system measurement characteristics, and performance logs. They also allow the user to select/deselect measurement functions.

Security Management SMFA

An organization's security management is composed of many diverse elements, most of which are outside the OSI scope. Security management SMFA provides the organization's security administrator with the facilities needed to support OSI security management, which is essentially the security of layer services and

OSI-Based Network Management

mechanisms. (The services and mechanisms themselves are defined in separate OSI standards.)

The security management SMFA standard defines the set of facilities that support security management; procedures associated with these facilities; the use of CMIS to provide the facilities; classes of managed objects in security; and subsets of facilities to allow conformance classes.

Every collection of open systems which intercommunicates has a common security policy. That security policy may be null, or it may be an intensive set of rules governing communication conditions. Furthermore, in any such collection of open systems, it is entirely possible for more than one security policy to exist in governing communications between particular open systems.

Typically, a security policy specifies how to protect against unauthorized reception or corruption of data transmissions between open systems. Also, the security policy specifies how entities on one open system may be granted access to resources on another open system. The security policy delineates how to determine the identity of entities wishing to communicate system to system—as well as the identity of their communications. The security policy spells out how (and to whom) to report significant security-related events as well as audit trail information.

A security policy's existence is defined in terms of OSI security services and mechanisms. Applications must provide some services and mechanisms that are not furnished by layer services and protocols. In either case, the security policy is implemented by selecting and configuring security services and by monitoring and controlling their operation.

Security management SMFA provides the set of facilities needed to operate on security objects, in order to implement the organization's security policy. Security management SMFA facilities allow management applications on different open systems to interoperate. Facilities provided or used by Security Management are:

- Security-Related Object Management;
- Security-Related Event and Audit Trail Management; and
- Security Management.

These three facilities allow operation of management applications on different open systems. These interacting applications manage security-related objects, attributes, states, and relationships via *security-*

related object management. (Security-related object management uses the same facilities as does the configuration management SMFA; namely, object configuration facility, attribute management facility, state management facility, and relationship management facility.)

Interactive management applications create and communicate security-related events, event logs, and audit trails via *security-related event and audit trail management.* (The facilities required for this function are similar to those defined in the fault management SMFA; namely, spontaneous event reporting, cumulative event gathering, event threshold alarms, event tracing, and management services contract.)

The third facility, *security management*, is only a placeholder in this SMFA. As of the May 1988 SMFA 4th Working Draft, the SC21 Working Group committee has not yet determined whether there is a need for facilities and primitives unique to a *security management* facility.

Accounting Management SMFA

Accounting management provides mechanisms to monitor information and control communications resources as they affect individual OSI users. This enables users and systems administrators to identify usage, as well as possible usage constraints. For those accounts in which usage incurs a cost, it may be recorded in cost units. As is true with other SMFA standards, the accounting standard relates only to the elements of an organization's accounting practices which are reasonably part of an OSI open system.

There are two aspects to accounting for OSI communications resource use: accounting within the communications medium itself and accounting for the use of communications resources within the end systems. The communicating end systems, and the communications medium itself, may indeed fall within different accounting domains. Each domain may exercise its own policy with respect to accounting management and report accounts for usage independently of the others. Thus, the ability to communicate accounting information from one domain to another is a necessary requirement.

There are a number of accounting procedures identified within both Application and Network Layers. These procedures allow users to activate accounts, collect accounting information, report accounting information, and negotiate for accounts. The accounting management SMFA standard covers those open systems management activities which imply the actual exchange of information within open systems.

OSI-Based Network Management

The standard specifies how CMIS and CMIP are used to effect accounting management and therefore defines the set of facilities and associated procedures to support accounting management, how CMIS services are used, the classes of managed objects involved, and subsets of the SMFA for conformance classes.

Specific terms defined within the standard include:

- Accounting System Administrator;
- Accounting Information;
- Accounting Unit;
- Accounting Limit;
- Unit of Charge;
- Accounting Report;
- Network Subscriber; and
- Application Subscriber.

Where accounting management activities are distributed among a number of computers, one system may require that the other systems exercise control over their communications usage through accounting mechanisms. Additionally, one system may require certain information from the other systems, such as data it can interpret to account for communications costs and data related to the other systems' communications costs. For example, System A may consolidate data from a number of other systems (B, C, D, etc.) to determine accountable usage. System A presents accounting summaries to network personnel or to systems B, C, D, etc. upon request. The accounting management SMFA standard sets no limit upon the reporting chain for accounting management data.

A system which generates accounting-related data is termed an *accounting management agent* (in our example, System A). An *accounting manager* is the term given to a system which requests or otherwise receives management data. Where an extended reporting chain exists, a system may perform the accounting manager role in respect to an agent further down the chain while also performing the agent role in respect to a manager further up the chain. There is no a priori requirement that managed objects which supply accounting data must be specific to the accounting SMFA. Any data made available by a managed object is deemed accounting data, if that data is obtained for accounting purposes.

The accounting management SMFA standard will eventually specify its facilities through CMIS/CMIP

standards. As of the June 14, 1988 Second Working Draft, however, the specifications are too preliminary to warrant review.

STRUCTURE OF MANAGEMENT INFORMATION (SMI)

The SMI standard is very important since it defines OSI elements used throughout all other OSI standards that are related to OSI Management. The SMI standard is also very abstract and is therefore somewhat difficult to fully appreciate.

The SMI standard defines the logical structure of OSI management information, which includes any information that may be the subject of OSI Management communications. This information is structured in terms of *managed objects*, their *attributes*, the *operations* performed on them, and the *notifications* that objects may issue. The SMI standard defines managed object concepts within the OSI information model and sets out the principles for *naming* the managed objects and their attributes. Objects must be named in order to be identified in OSI Management protocols.

The SMI standard also defines a number of *subobject types* and *attribute types* that are, in principle, applicable to all classes of managed objects. These definitions include the common semantics of the object/attribute types, the operations performed upon them, and the notifications they issue. The definitions also cover the relationships that may hold between the various types.

OSI elements defined in the SMI standard are listed in Table 1. The SMI standard defines the basic concepts of OSI Management information through several categories:

- Managed objects and their relationships;
- Attribute characteristics;
- Managed object characteristics;
- Management operations;
- Naming principles; and
- Object selection.

Directory

The Directory is a collection of open systems which cooperate to hold a logical database of information about a set of objects in the real world. The Directory

OSI-Based Network Management

provides capabilities required by OSI applications, OSI Management processes, other OSI entities, and the telecommunications services. Directory users, including both individuals and computer programs, can read or modify Directory information, provided they have the necessary authorization.

Information held in the Directory is collectively known as the Directory Information Base (DIB). The DIB is composed of Directory entities, each of which consists of a collection of information on one object. Each entity is made up of attributes, each with a type and one or more values. The types of attributes which are present in a particular entity are dependent on the class of object which the entity describes. This DIB information is typically used to facilitate communication between entities, users, terminals, computers, network elements, and miscellaneous OSI-related processing systems.

The Directory system plays a significant role in achieving a primary OSI goal. OSI's aim is to allow the interconnection of information processing systems which have very little in common. These processing systems have different manufacturers, are not always of the same age, are under the auspices of different managers, and have dissimilar levels of complexity. Yet, OSI is attempting to bring them together with a minimum of technical agreement outside the OSI standards themselves.

Among the capabilities which the Directory provides are those of "user friendly naming" (where objects carry names that are meaningful to average users) and "name-to-name mapping" (which allows dynamic binding between objects and their locations). Name-to-name mapping would enable self-configuring OSI networks, in the sense that adding, removing, or changing object locations would not affect OSI network operation.

The Directory provides a well-defined set of access capabilities known as the abstract service of the Directory. The abstract service provides a simple modification and retrieval capability. It is likely that the Directory will be distributed, perhaps widely so, along both functional and organizational lines. It has been

designed to support multiple applications, drawn from a wide range of possibilities. The nature of the applications supported will govern which objects are listed in the Directory and which kinds of access they will carry out. Applications may be very specific, such as the provision of distribution lists for electronic mail, or they may be completely generic in nature.

Definitions provided by the standard include:

- **Directory**—a collection of open systems cooperating to provide directory service;
- **Directory Information Base (DIB)**—the set of information managed by the Directory; and
- **User**—the entity or person which accesses the Directory.

The Directory standard provides implementation details within the following sections:

- Directory Information Base—DIB
- Directory Services
- The Distributed Directory
- Applying the Directory

Information within the DIB is arranged in the form of an inverted tree; that is, Directory entries are represented as elements in a directory information tree (DIT). Each entity in this DIT has a distinguished name, which unambiguously identifies that entry. Entries higher in the tree (nearer the root) will often represent objects such as countries or organizations, while entries lower in the tree will represent people, application processes, or other OSI-defined elements. The Directory standard defines a set of rules, the "Directory schema," to ensure the tree structure is properly maintained.

Directory services include interrogation of the Directory and modification of the Directory. Normal, error, and referred results are defined also. □

OSI Network Management/Forum: A Critical Assessment

This report will help you to:

- Assess the influence which the OSI NM/Forum holds over the future of network management standards and the NMS market.
 - Evaluate the progress of the OSI NM/Forum technical teams as they meet or near their stated goals.
 - Compare the strategic direction shared by OSI NM/Forum vendors with IBM's and Digital Equipment's.
 - Determine what impact, if any, current OSI NM/Forum activities will have on your future network management purchase decisions.
-

The OSI Network Management/Forum is a chartered consortium of 13 voting and 57 associate members, who are primarily vendors of network management equipment and/or carrier services. As it celebrates its first anniversary, the Forum proudly points to several key achievements which it accomplished on time or ahead of schedule. These milestones include publishing an OSI Protocol Specification and reaching a consensus on message sets for Fault and Configuration Management (officially called the "Application Services Specification").

According to the Forum's charter, the organization's goal is to supply specific implementation information, such as protocol options and message sets, that designers need to develop network management products and services that can work together. These specifications will be based on current Open Systems Interconnection (OSI) Management standards, to the extent they are defined.

Since its inception, many users (and some vendors) have been unclear on what the Forum really is and how it will go about accomplishing its stated goals. This confusion is augmented by the general lack of a clear understanding of OSI-based network management. Confusion, however, indicates interest—and more users are examining OSI as a serious long-term

network management direction. Consequently, users should attempt to understand what the Forum is and how it fits into the OSI picture.

RELATIONSHIP TO OFFICIAL STANDARDS DEVELOPMENT ORGANIZATIONS

The OSI NM/Forum is *not* a standards development organization—it has no formal association with the International Organization for Standardization (ISO), the International Electrotechnical Commission (IEC), the Comité Consultatif Internationale de Telegraphie et de Telephonie (CCITT), the American National Standards Institute (ANSI), or any other nationally or internationally recognized standards body. While the Forum's activities may enhance and improve the final outcome of ISO's OSI Management standards, the possibility exists that Forum influence may, in certain instances, be counterproductive to the global, long-term requirements addressed by standards development bodies. The risks and benefits of Forum activities are best understood when viewed within the context of the *standards life cycle*. Typically, standards evolve in three stages: development, implementation,

OSI Network Management/Forum: A Critical Assessment

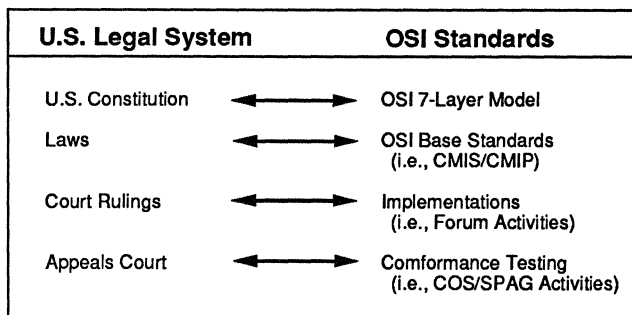


Figure 1. The relationship between standards development, standards implementation, and standards conformance testing is analogous to the structure of the United States government.

and testing. Each phase is distinct, and the related work is accomplished via a different type of organization. (See Figure 1.)

Standards Development: ISO, CCITT, and Other Standards Development Organizations

ISO/IEC Joint Technical Committee 1 (JTC1) is responsible for *developing* OSI Management standards, the most familiar of which is the pair of Draft International Standards (DIS) 9595 and 9596, which specify the Common Management Information Service/Common Management Information Protocol (CMIS/CMIP). As a *standards development* body, ISO/IEC/JTC1 strives to maintain a global, long-term perspective as it defines base standards which must be applicable to varying technologies (such as voice and data), available in all countries. To achieve such universality, the development phase is, by nature, very conceptual and a technically and politically complex process. Thus, standards developers realize that the resulting definitions are likely to contain some omissions and ambiguities. To help bridge gaps, and to achieve the greatest degree of universal applicability, standards developers present a mix of options, from which vendors/users may choose to create real products.

Standards Implementation: OSI NM/Forum, MAP/TOP, and the NIST SIG

The primary task of the implementation phase is for vendors and/or other interested parties take the base standard, compare it to their requirements, and choose the appropriate options to meet their needs. During this process, implementation groups should identify the base standards' ambiguities and omissions, and relate them (as well as proposed solutions) to the standards development bodies. For example, Manufacturing Automation Protocol (MAP) implementation specifications address resource initialization and termination to a greater extent than do OSI

standards. In particular, MAP implementation guidelines specifically address the requirements of an 802.4 factory environment, while OSI must remain totally implementation dependent and apply to more diverse settings, including wide area and local area networks.

Standards implementers provide a reality check on defined standards. Implementers create industry pressure, which can act as an indirect catalyst for pushing an appropriate subset of standards through ISO committees; without this positive pressure, standards bodies are likely to wade through the entire set of global requirements before producing anything final. In essence, the Forum has the potential to push standards developers to produce a "phased approach," and approve a critical minimum subset of network management standards within the next 18 to 24 months.

The temptation which the OSI NM/Forum member companies and other implementers face is, in the name of expediency, to fix the holes in the standard within the implementation group instead of making recommendations to standards developers for their resolution. The total global set of requirements addressed in the standards bodies is beyond the mission of those outside the standards development process. Fixing the standard in the implementation phase creates expensive problems—forcing vendors to decide between implementing the interim fix or waiting until the fix is modified and made permanent within the standards development groups. The expense is incurred when some vendors choose the interim while others choose the final version—creating an expensive chasm. As always, this expense will be borne by the user.

The cost of multiple implementations could be even greater without the Forum, however. For the past two years, vendors have felt increasing market pressure to produce a multivendor network management solution. Without a consortium such as the Forum, some vendors would probably attempt to implement OSI-based network management on their own before standards were finalized. This would produce a number of incompatible solutions, which would be virtually worthless.

When done the right way, both developers and implementers can take advantage of the expertise and progress made in both groups, and can feed into the good will of both. Often, as in this case, key vendors have representatives on both standards bodies and implementation groups. This is not only expedient but a necessary part of developing consensus and team spirit among member companies whose economic interests may be at odds with each other. Vendors that can cooperate within the standards development phase, and operate as a team within those circumstances, and in

OSI Network Management/Forum: A Critical Assessment

AT&T and the OSI NM/FORUM

Although not the sole instigator, AT&T was certainly a major force in founding the Forum. AT&T was apparently dissatisfied with the pace of progress among standards development groups, and needed to generate vendor support for its Network Management Protocol (NMP)—an implementation of OSI CMIS/CMIP. AT&T is currently developing OSI-based network management interfaces for all of its network services. AT&T plans to introduce an enhanced network service product featuring an NMP interface before the end of 1989.

In addition to meeting OSI requirements for its other products, AT&T had to boost support for NMP to make its ACCUMASTER Integrator product line viable. While Forum members who based their work on CMIS/CMIP documents are under no constraint to accept AT&T's dictum, their resulting Protocol Specification is very similar to NMP.

While Forum activities may increase momentum behind support for AT&T's NMP and its Unified Network Management Architecture (UNMA) in the short term, it will soon fade if AT&T's ACCUMASTER Integrator product fails to live up to expectations. In other words, helping to found the Forum was a smart move, and it brought AT&T's products a lot of publicity, but it may end up backfiring on AT&T if the Integrator is not a stunning success.

Even if users become disillusioned with AT&T's Integrator and UNMA, other Forum members will continue to benefit from joining in network management standards implementation. It is likely that key NMS equipment members (particularly Hewlett-Packard, Unisys, Amdahl, and DCA) will forge something useful out of the process, even though the Forum may not have been their idea to begin with. Many Forum members are truly committed to OSI-based network management, and despite OSI's omissions and its telecommunications hardware bias, this year's implementation efforts will produce a more informed, experienced vendor community.

On the other hand, carriers, including AT&T, may end up losing ground because users are now focusing more on OSI than on ISDN. Although the two standards are closely intertwined (and ISDN services will eventually require OSI-based network management), carriers may regret not putting more resources into promoting ISDN and other services that are more directly within their realm of expertise. Users today don't look at the PBX as a LAN integrator, as AT&T hoped; similarly, the users of tomorrow may not look at carrier-provided INMSs for integration either, despite the excessive time and money which AT&T and other carriers are spending on convincing the MIS world that they can provide an integrated network management solution for data—and voice—networks.

the implementation phase, will have the best momentum to successfully and profitably implement the standards.

CONFORMANCE TESTING PHASE: COS/SPAG

Within the U.S., the Corporation for Open Systems International (COS) is the chief vendor/user consortium working to establish conformance testing of multivendor OSI and ISDN products. COS now coordinates its efforts with the Standards Promotion and Application Group (SPAG) in Europe, and is seeking to work with the Promotion of OSI (POSI) in Japan.

COS/SPAG's scope is broader than the Forum's, however, as it encompasses all of OSI—not just OSI-based network management. To sharpen their focus on network management, COS/SPAG and the Forum have created the Forum/COS/SPAG Executive Council. This group is chartered to work on common conformance testing goals; one of the council's first tasks is to establish conformance testing procedures for products demonstrated at Showcase '90, scheduled to begin in the fall of next year. (See "USERS: WHEN CAN

THEY EXPECT TO SEE OSI NM/FORUM PRODUCTS" for more information on Showcase '90.)

COS/SPAG and the Forum are currently evaluating alternatives to support conformance testing for Showcase '90 and beyond. Together, the three groups face the task of resolving a number of key issues—the question of "how do you test Network Management" has yet to be settled. Open technical issues include establishing the location of "testers" within the protocol stack, as well as determining the system behaviors to be tested. At this point, over 200 such "behaviors" have been identified for CMIP.

In addition to checking protocol compliance, it is undecided whether it is appropriate to test the "real effects" of management commands on actual network devices. For example, if an NMS issues a message to turn off the modem, to test the real effect, you must go out into the network and determine whether the modem was indeed turned off. At this point, conformance testing groups have not decided whether it is appropri-

OSI Network Management/Forum: A Critical Assessment

VOTING MEMBERS:	ASSOCIATE MEMBERS: (Continued)
Amdahl Corp.	ING. C. Olivettie, and C. S.p.A.
AT&T	Interlan, Inc.
British Telecom	Kokusai Denshin Denwa Co., Ltd.
Digital Communications Associates, Inc.	McDonnell Douglas Network Systems Company (Tymnet)
GEC Plessey Telecommunications, Ltd.	Microtel
Hewlett-Packard	NCR Corp.
MCI Telecommunications	NEC America, Inc.
Nippon Telegraph and Telephone	Netlabs
Northern Telecom, Inc.	Network Equipment Technologies
Societa Finanziaria Telefonica S.p.A. (STET)	Newbridge Networks Corp.
STC plc	Nixdorf Computer Engineering Corp.
Telecom Canada	Novell, Inc.
Unisys Networks (Timeplex)	OKI Electric Industry, Co., Ltd.
ASSOCIATE MEMBERS:	Paradyne Corp.
Alcatel, n.v.	Philips Data Systems
Applied Computing Devices, Inc.	Prime Computer, Inc.
Atlantic Research Corp.	Protocols Standards and Communications (PSC), Inc.
Avant-Garde Computing, Inc.	Racal-Milgo
Bull SA	Racal-Milgo Ltd.
Cable and Wireless plc	Retix
Case Communications, Ltd.	Siemens AG
CNCP Telecommunications	Sirti, S.p.A.
Computrol	Spider Systems Ltd.
Contel Technology Center	Stratus Computer, Inc.
Data General Corp.	Synoptics Communications, Inc.
Dynatech Communications	Systems Reliability plc
Ericsson Business Communications AB	Tandem Computers
France Telecom	Tech-Nel Data Products, Ltd.
Fujitsu America, Inc.	Teknekron Communications Systems, Inc.
Gandalf Data Ltd.	Televerket Sweden
Gartner Group	Telenet Communications Corp.
General Datacomm, Inc.	Telindus n.v.
Hekimian Laboratories, Inc.	Telwatch
Hitachi Telecom (USA), Inc.	Ungermann-Bass, Inc.
Infotron Systems Corp.	Vance Systems, Inc.
	Zellweger Telecommunications

Table 1. OSI Network Management/Forum members, as of July 15, 1989.

ate for NM to test real effects, since it pertains more to the network than to the interaction between network management systems.

OSI NM/FORUM: WHY IT WAS FOUNDED AND ITS INFLUENCE ON NMS MARKET

While AT&T was a major force in founding the Forum, the other voting members had significant input in creating the group's charter and organizational structure. The Forum's resulting mission and planned accomplishments reach beyond AT&T's initial motivations (see first sidebar). In particular, British Telecom and STC Plessey developers had important roles in forging pre-Forum documents which outlined the status of OSI Management standards and what needed to be accomplished to bring OSI-based network management to the market in a timely fashion. Early discussions among all eight founding Forum vendors produced a consensus that only a vendor consortium could accelerate the birth of an already overdue concept—multivendor network management.

The antithetical force acting as a catalyst to the Forum's founding (and also, in part, to the beginnings of the entire OSI movement) is IBM and its Systems Network Architecture (SNA). OSI represents the only viable solution for users who wish to create large-scale networks without being tied to a single-vendor solution, i.e., SNA—and the only means by which other vendors can afford to offer multivendor network management products. OSI represents a non-IBM choice. How timely, practical, and economical that choice will be is an important question, but outside the scope of this report.

OSI/NM Forum vendors, including AT&T, must develop a non-SNA approach for network management—and face the dilemma of giving their customers a choice without breaking vendor R&D budgets. Banding together was a logical thing to do, and in light of the Forum's membership growth over the past 12 months, a lot of vendors must agree.

How will the next year's Forum activities influence the NMS market and the network management platforms

OSI Network Management/Forum: A Critical Assessment

Board of Trustees:	
Title	Company/Name
Chairman	Amdahl Corporation Leighs Church, Director, Software Product Management
President	AT&T John Miller, Director, Network Mgt Market Planning
Vice President	British Telecom Keith Willets, Manager, Communication Management Strategy
Secretary	Northern Telecom Ian Sugarbroad, Director, Business Development
Treasurer	Telecom Canada Robert Montgomery, Director of Business Development
Trustee	STC/Telecom Systems Andrew Roberts, Marketing Director
Trustee	Hewlett-Packard David Mahier, Product Marketing Manager, Information Networks Division
Trustee	Unisys Networks Larry Thomas, Manager Advanced Systems Architecture
Trustee	DCA Mani Subramanian, Vice President of Development

Table 2. OSI Network Management/Forum officers.

of participating vendors? If the Forum stages successful interoperability demonstrations in autumn 1990, OSI-based network management will indeed appear closer to reality. While Forum demos do not directly influence standards development, they will create an impression upon the industry that OSI Management is progressing—many users will incorporate “OSI compatibility” not only into their strategic direction but also into their implementation plans as well. This could be a positive influence; just how substantive it will be is not yet known.

Perhaps more important is the Forum’s potential impact on persuading ISO/IEC/JTC1 to approve a critical minimum subset of network management standards in the short term, as mentioned previously. This phased approach could open the door to truly OSI-compliant network management

implementations—which pose less risk than incompatible interim implementations.

PROGRESS REPORT ON THE FORUM’S TECHNICAL TEAMS

The Forum has produced two significant documents in the past year: the Protocol Specification and the Application Services Specification. Producing these documents was no easy task; however, the real tests lie ahead—producing a meaningful Object Specification and staging interoperability demonstrations. Beyond that, the Forum still must address issues beyond the scope of Fault and Configuration Management (the essence of the Application Services Specification) if it is to produce a lasting effect on the future of network management.

To put the Forum’s achievements and goals into perspective, it is helpful to first review the Forum’s structure and meeting schedule.

THE BALANCE OF POWER

The Forum comprises over 70 members; 13 vendors are voting members (see Table 1). For the first year, voting members pledged \$40,000 per year and two engineering years (dues for the second year are about to be increased). Voting members are entitled to participate in all policy and technical activities. The Board of Trustees, which possesses the ultimate authority in making nonvoting decisions, comprises representatives from voting member companies (see Table 2).

Voting members also form the core of the five technical teams; associate members participate in the technical teams by request. Associate members pay \$5,000 per year to participate in general meetings and to receive all published Forum documents.

The five technical teams are each led by a Team Leader whose primary focus is to bring the group to a timely resolution of the matters at hand. Each team has one or more Editors, who have a pivotal role in creating the technical document and maintaining it on behalf of the team. Most teams also have Subject Matter Experts who are responsible for insuring technical accuracy. Each Technical Team Leader reports to the Technical Director, who is a member of the Board of Trustees. During the Forum’s first year, the Technical Director was a British Telecom representative; the post will soon be held by a representative from DCA. The five technical teams are:

- **Protocols (P-Team):** This committee has already completed its most important goal: the publication of the Protocol Specification. Hewlett-Packard

OSI Network Management/Forum: A Critical Assessment

IBM AND DIGITAL: LOOKING IN FROM THE OUTSIDE

IBM and Digital Equipment were the two obvious omissions on the Forum's list of founders. The Forum invited both vendors to join, and throughout its first year of existence, maintained an ongoing dialog with each pertaining to their questions and concerns. Both IBM and Digital, however, had long-range plans of their own, developed long before the Forum appeared on the scene.

IBM

Unlike most OSI NM/Forum vendors, IBM has had a comprehensive NM strategy for several years now. That strategy must work *now* for the thousands of IBM users dependent upon the company's present equipment and software. It is impossible and unreasonable for IBM to abandon NetView or SNA in any fashion. NetView is burdened by its association with outdated technology (VSAM instead of a relational, object-oriented database, 3270s instead of color-graphics workstations). Even so, IBM is working hard to catch up to the future and evolve NetView to support OSI—just as the company is evolving SNA into a more peer-to-peer structure.

While IBM may be slower to implement "leading edge" network management technologies, such as OSI, the company has its ear to the ground and knows what its customers need *today*. It will take at least two years before OSI-based network management systems will come close to providing the functions for **multivendor networks** that NetView delivers right now for SNA networks.

IBM understands network management and straightforwardly addresses the complexities introduced by application errors and human errors, which typically occur more frequently and prompt more trouble calls than hardware faults. IBM solidly addresses Data Center Management (which it calls "Systems Management.") Although OSI standards can encompass Data Center Management, network concerns (such as telecommunications hardware and circuits) are currently receiving more attention within standards development and implementation groups.

Although IBM well understands network management and still dominates the NMS market, it recognizes the swelling OSI tide and is slowly moving to accommodate it. Though OSI may be at odds with SNA in a conceptual sense, IBM must eventually make them work together under NetView.

IBM previously stated that it would not join until the Forum addressed several critical issues, including the Forum's effect on the ongoing efforts of COS and SPAG. IBM is active in COS/SPAG and has numerous representatives on ISO/IEC/JTC1 committees. The Forum has addressed IBM's concerns by forming the Forum/COS/SPAG Executive Council.

IBM has also stated that it prefers to wait until ISO/IEC/JTC1 standards have been finalized before embarking on OSI-based network management implementation. This interim "waiting period" may give IBM enough time to piece together a complex OSI/SNA puzzle—one which *must work* for thousands of IBM customers. Accelerating the OSI movement—the Forum's mission—would only shorten that interim.

IBM now sees the handwriting on the wall, however, and is about to become privy to Forum activities by joining up. This has a positive effect on the Forum, and hopefully will help to increase attention to the critical needs of managing Data Centers in a multivendor environment.

DIGITAL EQUIPMENT CORPORATION

Digital has already invested almost three years in its network management platform, called Enterprise Management Architecture (EMA). During this time, Digital has been active on ISO/IEC/JTC1 committees responsible for defining CMIS/CMIP, the Management Information Base (MIB), and the Structure of Management Information (SMI).

Digital is structuring EMA to accommodate OSI Management objectives in an attempt to ease the transition from Digital's proprietary protocols (such as Network Information Control and Exchange protocol, known as NICE) to CMIS/CMIP. For example, EMA includes one centralized Management Information Repository (MIR), which essentially performs the same role as the OSI Management Information Base (MIB). In this sense, EMA is much more amenable to OSI than IBM's platform, which is evolving now to include two NM databases (an Enterprise and Operations database). Likewise, Digital designers are using OSI concepts and terminology in their architectural models. How far Digital's interpretation of those concepts differs from ISO/IEC/JTC1's intent will not be known

OSI Network Management/Forum: A Critical Assessment

IBM AND DIGITAL: LOOKING IN FROM THE OUTSIDE (Continued)

until EMA's actual implementation—which is at least two years away. While the Forum seems intent on accelerating OSI, Digital is attempting to set realistic expectations and taking the time it deems necessary to do the job right.

Despite Digital's initial reluctance, it is now planning to join the Forum. The company has closely watched the group's activities for over a year now, and assessed both the Forum's direction and marketplace pressures. In the interim, however, Digital aligned its own mini-consortium of seven vendors who are committed to jointly hammering out EMA implementation details.

Some of these vendors, including DCA, are also key Forum members.

EMA's initial release is expected to support proprietary protocols, such as NICE, as well as CMIS/CMIP. During its implementation efforts, Digital is undoubtedly uncovering omissions and ambiguities in the base standard. Up to this point, Digital has been developing its own solutions to these problems (in conjunction with the other seven EMA-supporting vendors). How these solutions may differ from the OSI NM/Forum's is a matter of speculation. When Digital joins the Forum, it will not have to guess.

chaired the committee and, after a successful term, has passed the baton to Northern Telecom. An AT&T representative served as Editor for the original Protocol Specification document.

- **Messages (M-Team):** AT&T chairs this team, which has just completed its first major milestone—publication of the Application Services Specification, which outlines the message sets for Fault and Configuration Management. British Telecom provided the Editor for the technical document.
- **Object Definitions (O-Team):** This is where the action is now and where some of the most difficult tasks remain to be tackled. All voting members have representatives on this team. The O-Team was led by Northern Telecom during the first year, and is now under the direction of an Amdahl representative, who previously served in an editorial capacity. AT&T also has an Editor on the O-Team.

One "gap" in current OSI Management standards is how managed objects (such as modems, lines, terminals) will be defined. In the initial Object Specification, the O-Team will define a set of about 20 objects needed by Forum members for initial implementation (particularly for Showcase '90). Subsequently, the team will describe procedures for additions and extensions to that set.

Perhaps more critical than choosing the object set is defining the object *template*, upon which object definitions are based. Attributes, behaviors, actions, name formats, etc. make up the template or *schema*. The Forum's O-Team is now developing its own templates and what they will contain, using ISO/IEC/JTC1 working documents as a starting point. The team hopes to publish its Object Specification by the end of this year. If the Forum succeeds in creating a viable Object Specification, non-Forum vendors may lag behind Forum members in implementing

OSI-based network management (see second sidebar). If ANSI (as the U.S. member of ISO/IEC/JTC1) presents a viable alternative to the Forum's object templates during the next 18 months, there may be a tug-of-war between ISO and the Forum for some time. Many Forum members are working hard to avoid that embarrassment, however. No matter what the outcome, object definition is a never-ending cycle, and this committee's work will probably continue for as long as the Forum remains in existence.

- **Architecture and Services (A & S-Team):** This group must determine such long-range issues as how one management domain relates to another, and how all Forum-approved messages and protocols fit into the overall scheme of network management. This committee began its work in March 1989, is currently chaired by a representative from British Telecom, and has several Editors.
- **Testing Team:** This committee's immediate focus is to establish testing procedures for the Showcase '90 demonstrations, which will comprise multiple, separate events in different cities from September 1990 through January 1991. (See the section entitled "USERS: WHEN CAN THEY EXPECT TO SEE OSI NM/FORUM PRODUCTS?")

There are several key committees in addition to these technical teams: the Strategy and Planning Committee (led by an STC representative) responsible for determining future work priorities, including the agenda after Showcase '90; the Finance and Audit Committee, led by Telecom Canada; the Membership and Nominating Committee, led by Northern Telecom; and the Publicity and Events Committee, led by Amdahl, which is working on the logistics of pulling together the Showcase '90 interoperability demonstrations.

OSI Network Management/Forum: A Critical Assessment

MEETING STRUCTURE

The Forum's Board of Trustees meets every six weeks, as do many of the committees listed in the preceding paragraph. There are three general meetings a year which both Voting and Associate members may attend. The last meeting was held in June in London, as was the Forum's first Annual Meeting.

Technical teams meet together approximately every six weeks, and plenary sessions are held periodically, when all five technical teams meet together in one location. The first technical plenary session was held in July 1989.

ASSESSING PROGRESS TO DATE

Although its members are committed to common goals, the Forum is composed of competitors—the group is, by nature, in conflict with itself. Despite this, the Forum survived all of the political maneuvering during its first year and managed to produce two important documents (Protocol Specification and Application Services Specification) on schedule. If that trend continues, the group may achieve a workable Object Definition—but the complexity and critical nature of that task will test the Forum's team spirit, the expertise of its members, and its commitment to a larger common goal.

USERS: WHEN CAN THEY EXPECT TO SEE OSI NM/FORUM PRODUCTS

The Forum has scheduled a number of different interoperability implementations during the fall of 1990. Dubbed "Showcase '90," these events will be staged at

different locations in both the U.S. and Europe. Showcase '90 will not be the same implementation in each city, however, and no one vendor will participate in all six events. Rather, the Forum envisions having a "cluster" of two or more vendors work on staging an interoperability implementation in one particular city. It is likely that vendors with standing joint agreements (such as DCA and Northern Telecom) will align themselves into a cluster, although new agreements may be forged as a result of planning for Showcase '90.

If successful, Showcase '90 will make a major stride in establishing the Forum's credibility. Even so, Forum vendors will still be a long way from delivering true multivendor network management to users. While exchanging alerts between two or three vendors' products is a milestone, it does little to help the user who has equipment from a different mix of vendors. More importantly, we have all seen product demos at trade shows, and waited months (or more) for real products to appear on the market. While the Forum pledges to showcase only "real" products, it is unrealistic to expect that comprehensive Forum-approved products will be right around the corner.

Showcase '90 will expedite the implementation phase of OSI-based network management standards, however, and surely the participating vendors will benefit from the experience. Again, if those vendors transmit the knowledge they therefore gain back to the standards development committees, both users and vendors will profit. If Forum vendors attempt to force interim solutions onto the market before the standards can be amended to accommodate their input, early products will have to migrate to the eventual standards—perpetrating dichotomy and confusion for some time to come. □

Toward a Unified Theory of Managing Large Networks

This report will help you to:

- Manage integrated networks supplied by different vendors.
 - Use emerging standards for integrated network management.
-

The Internet—the network connecting 60,000 computers at universities, government facilities, and think tanks across the United States—came to a standstill last November. What was described as a computer virus (although strictly speaking it was a worm, able to replicate on its own) spread through the network like an electronic wildfire. The organizations that escaped the worm were those that immediately cut their ties to the Internet or that earlier had adopted measures to stop such security breaches. But even managers on those portions of the network that were able to immunize themselves had only a partial idea of what was happening.

They too lacked the measures to determine whether security had been violated and to figure out the best approach to isolating affected portions of the network. Many organizations reverted to their primary management tool: a frantic telephone call to whom-ever was overseeing the network. The incident, which received nationwide publicity, forced some organizations to take a closer look at network management issues.

A more systematic management approach marks a change from when networks were just a laboratory curiosity or a showpiece for affluent organizations. Since AT&T Co.'s 1984 divestiture, and with the growing popularity of local-area networks (LANs), much networking equipment has come to market.

This Datapro report is based on "Toward a Unified Theory of Managing Large Networks," by Paul J. Brusil, The Mitre Corp., and Daniel P. Stokesberry, National Institute of Standards and Technology. © 1989 by IEEE. Reprinted, with permission, from the *IEEE Spectrum*, April 1989.

Familiarity with data communications gives scant protection against difficulties encountered in stringing together equipment from many vendors. Mitre Corp.'s network supports some 2500 users at several locations in Massachusetts and 1000 more at sites near Washington, D.C. (The company advises the U.S. government on technology acquisition.) It has been involved with computer networking for more than 20 years and helped develop protocols used by the ubiquitous Ethernet LANs. But it still has its own problems managing networks whose list of equipment resembles a catalogue from a data communications supply house.

A user seated at an IBM PC at Mitre's headquarters in Bedford, Mass., may request that a file be sent to other offices a short way down the road. The personal computer, equipped with a communications card from Excelan Corp., San Jose, Calif., is tied to an Ethernet connecting a few users in a workgroup. The LAN will then send the file to equipment from Aplitek Corp., Wakefield, Mass., or Chipcom Corp., Waltham, Mass., that modulates the baseband for transmission over a broadband network.

From there, the file hops to a bridge from Vitalink Co., Fremont, Calif., that connects the broadband link to a digital T1 line—one running at 1.544 megabits per second—that ties headquarters to satellite offices. The whole process then continues until the data package reaches the user for which it is intended. (Mention of any products in this report does not constitute an endorsement by the authors or their organizations.)

Many links and boxes on the network have their own sets of configuration and diagnostic tools, each with a

Toward a Unified Theory of Managing Large Networks

different command language and different ways of representing data. Some others, however, have no such tools, while some of those that do have them furnish row upon row of statistics that must be deciphered. Others supply just one or two numbers that give little idea about what is really going on. Mitre's managers and technicians must run what amounts to a mini-utility, integrating all these systems into a cohesive service.

If they could instantly change things, those in charge would opt for a standard network management framework with standard data formats and protocols to bring this flood of management information together. With hindsight, standards groups are now addressing the issue of network management.

WHAT IT TAKES TO MANAGE

Control and monitoring of the network are broken down by standards groups into what are called system management functional areas (SMFAs) that describe what requirements management systems fulfill. The five areas are:

1. Configuration: initiating or terminating the system's operation; obtaining status information—whether it is still working; and mapping the network's links and nodes.
2. Performance management: gathering information about network performance, such as measurement of throughput, availability, errors, and reserve, all needed for performing maintenance or planning network growth.
3. Fault management: detecting and correcting problems in disabled components, including setting alert parameters and performing diagnostics when a node or link fails.
4. Security: maintaining passwords and encrypted data links, and keeping records of possible security violations.
5. Accounting: determining which individuals or groups use the network, and keeping track of their billing and payments.

An SMFA constitutes an idealized model of how network management should work. With the exception of the telephone companies and other providers of commercial services, few private networks have comprehensive management schemes in place. Where control mechanisms do exist, they tend to be in the areas of configuration and fault management. Many

organizations hope areas targeted for standardization may provide models for the planning of network management.

Lacking standards, however, network managers must find makeshift solutions. A few have established their own de facto standards and buy most of their equipment from one systems vendor, who also supplies proprietary protocols and databases for exchanging and storing management information.

Most organizations, however, have equipment too varied for such a simple approach, and some vendors prefer not to link their equipment to a competitor's network management system. As a result, the number of consoles in a network control room may make it look like a space-launch command center. In other networks, the management systems are dispersed throughout the organization and trying to fix a problem may require a surfeit of three- and four-way conference calls.

Inadequate technical tools means that network personnel live in a state of chronic crisis. Without standards, Mitre and other organizations have had to act as their own systems integrators. An application called NetMan—20,000 lines of code for a Sun 4 workstation—displays status information about each vendor's network-management system on the one screen.

Even with NetMan, however, Mitre's network managers must be polyglots who have mastered the intricacies of a separate command language and display formats for each workstation window. Although comprehensive tools for collecting statistics are yet to come, the current approach lays the groundwork for further integration as management tools improve. This is where standards will help.

STANDARD-BEARERS UNITE

Mitre and other users are looking to international standards for management integration. Standard making for network management is rooted in efforts by the International Organization for Standardization (ISO) to set in stone its Basic Reference Model for Open Systems Interconnection (OSI). This seven-layer model, published in 1984, is not restricted to network management; it provides a common basis for the coordination of standards for establishing and managing communications.

Standards groups have used the OSI model to develop protocols for specific technologies: LANs, long-distance data networks (wide-area networks), networks for connecting up a city (metropolitan-area

Toward a Unified Theory of Managing Large Networks

<p>Common Management Information Protocol (CMIP): specification of the bit patterns needed to transmit a request for a network management service.</p>
<p>Common Management Information Service (CMIS): building blocks needed to carry out a network control or monitoring function, for example: SET, GET, or EVENT REPORT.</p>
<p>Managing and agent processes: a management process initiates network management monitoring and control transactions; an agent process receives these transactions and may also send an EVENT REPORT.</p>
<p>Systems Management Application Entity (SMAE): abstract description of the communication processes for performing a management function.</p>

Table 1. Terms defined.

networks); methods for connecting networks (gateways, bridges, routers); and protocols that specify how distributed processing systems establish physical links or how applications running on such systems communicate. There are seven protocol layers that supply these services: physical, link, network, transport, session, presentation, and application.

The OSI Model breaks down network management into systems management and layer management. Systems management uses all seven OSI layers for monitoring and controlling the network. Layer management acts directly at a single layer.

So far, most ISO work on management has been dedicated to systems management standards. Here, a managing process communicates with its peer, or agent process, in a remote part of the system. The standards classify interaction between manager and agent as a set of operations related to the basic tasks of network management, such as configuration.

Manager and agent processes reside at OSI's application layer, which provides a software interface for application programs that are to use the network. (See Table 1.) ISO calls the management communication parts of these processes a Systems Management Application Entity (SMAE) (see Figure 1). Like any other OSI application, each SMAE takes advantage of the underlying protocols supplied by OSI: presentation, session, transport and so on.

One component furnished by SMAE is a set of service primitives called Common Management Information Service (CMIS). Primitives are building blocks that define actions needed to perform a function. The main services are: INITIALIZE/TERMINATE—establishing a link between a managing and agent process;

GET—obtaining information from an agent process; SET—setting a variable; and EVENT REPORT—sending a report of the number of errors or some other network event.

The technical description of the protocol elements to provide CMIS is specified in another standard, the Common Management Information Protocol (CMIP), which details the structure and encoding of the protocols.

Exchanging information between SMAEs, often implemented as a process within a computer's operating system, is only part of what manager and agent processes do. The agent process must reach out to the managed object, which is a physical or logical system. A managed object can, among other things, be a protocol, a communication link, or a physical device such as a modem. The Structure of Management In-

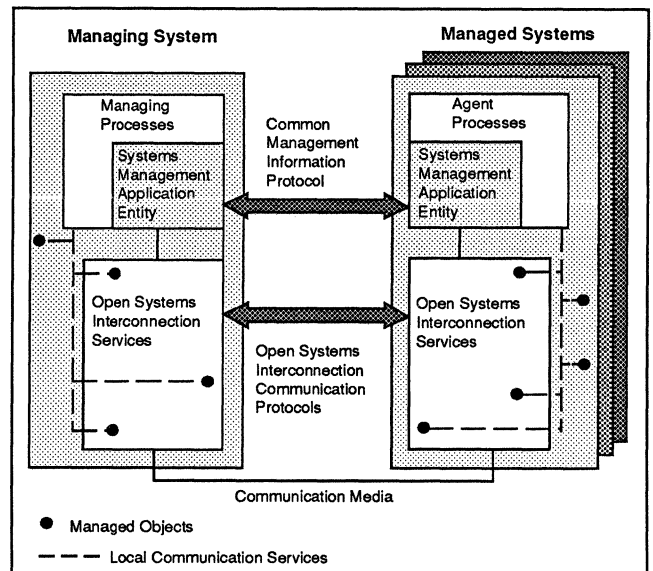


Figure 1. A managing process, which can initiate a network management transaction, takes advantage of the underlying services of the Open Systems Interconnection standards' layers to relay a management command to a remote agent process. The Systems Management Application Entity (SMAE) uses Common Management Information Protocol to exchange information with the remote agent process. The SMAE exists as a process in an end system, such as a computer. The final step requires that an operating system or other proprietary resource of this end system be used to complete the transaction with a managed object—a modem, a protocol, a subnetwork, or any other system. The managed objects within the OSI services box (protocol timers or packet counters, for instance) are at different protocol layers that correspond to the OSI Basic Reference Model, which provides a framework for standard communications. Those outside the services box represent objects managed within a vendor's proprietary system, such as a computer's serial number or password tables. The physical connection between two remote systems takes place over communications medias such as a local-area network.

Toward a Unified Theory of Managing Large Networks

Management Area	Standards Organization ¹	Work by Standards Groups	Status ²	Estimated Completion Date
Architecture	ISO SC21/WG4 ISO 802.1 CCITT SG VII	OSI management architecture	IS	Complete
		LAN layer-management architecture	WD	Undecided
Management communication services and protocols	ISO SC21/WG4 CCITT SG VII IAB NetMan	Telephony network-management architecture	Work starting	1990
		Common Management Information Services (CMIS)—to carry out a network control and monitoring function; Common Management Information Protocol (CMIP)—bit patterns to transmit a request for a network management service	DIS Work starting RFC	1989-1990 1990 Complete
		LAN layer-management protocol	WD	Undecided
		Simple Network Management Protocol, a transition protocol for managing the Internet before OSI's CMIS and CMIP are deployed	RFC	Complete
System management functions	ISO SC21/WG4 CCITT SG VII, and ANSI T1M1.5	Configuration and fault management	WD	Undecided
		Performance, accounting, and security management	WD	1991-1993
		Common functions—state management or error reporting used in systems management, for example	DP	1991
Managed objects	ISO SC21/WG4	Defining structures, formats, and guidelines for managed object definitions (Structure of Management Information)	DP	1991
	ISO SC21/WG4 and WG5	Defining parameters to be managed for systems (WG4: systems identification and serial numbers, for example) and ISO upper-layer protocols (WG5: which system is to initiate sending, for example)	Ranges from work starting to DP	Undecided
	ISO SC6/WG2 and ISO SC6/WG4	Defining parameters to be managed for ISO lower-layer protocols (timers specifying retransmission timeouts and counters registering number of packets sent, for example)	WD	1991
	IEEE 802.2-802.10	Defining parameters to be managed for lower-layer protocols for LANs and metropolitan-area networks; includes security	Ranges from beginning efforts to DIS	Undecided
	ANSI ASC X3T9.5	Defining parameters to be managed for high-speed, fiber-optic LANs (fiber distributed-data interface)	Work starting	Undecided
	ANSI ASC T1M1.5	Defining parameters to be managed for telecommunication devices such as multiplexers	WD	Undecided
	CCITT, various SGs	Defining parameters to be used in communications such as those for Integrated Services Digital Network, a standard network for combining voice, data, and other digital services	Work starting	Undecided
	IAB MIB WG	Defining parameters to be managed for the Internet's Transmission Control Protocol/Internet Protocol	RFC	Version 1 complete

¹ISO: International Organization for Standardization; SC: subcommittee; WG: working group; CCITT: International Telegraph and Telephone Consultative Committee; SG: study group; IAB: Internet Activities Board, which oversees the Internet; ANSI: American National Standards Institute; ASC: Accredited Standards Committee; MIB: Management Information Base.

²IS: International standard; WD: working draft; DIS: draft international standard; RFC: request for comment (IAB's equivalent of a standard); DP: draft proposal. ISO standards work proceeds from working draft through draft proposal to draft international standard to international standard.

Table 2. Network management standards.

formation (SMI) standard describes formats and structures needed to define the attributes of each class of managed object: protocol timing thresholds for retransmitting a packet are examples of an attribute.

Standards groups produce specifications for managed objects, using SMI, which will be used to assemble a standards library, called a management information base.

Toward a Unified Theory of Managing Large Networks

OSI management, however, does not dictate how to communicate between the agent's SMAE and the managed objects. This is the interface between the standard part of the network and proprietary systems. After recognizing a CMIS service request (SET or GET, for instance), the agent process—in conjunction with the agent's operating system—translates the managed object information into proprietary data and command structures, which may be implemented, say, in the firmware of a communications card plugged into a personal computer's backplane.

OSI scrupulously avoids dictating the final design of network management software. Among decisions not part of standards that a vendor must face are how quickly management information is displayed, and the graphics to be used on a control console.

Layer management is another form of network management for which standards are under development. A LAN might be a candidate for a layer-management standard. It provides a physical path for data to move from one computer to another, so that LAN communication standards conform to OSI's two lowest layers, physical and link. A LAN layer-management standard, for example, might manage the flow of packets through a bridge connecting two LANs at the link layer.

The emerging standards will be of little value, however, if no one uses them. As with other protocol layers, groups of vendors have agreed to implement standards in actual products. One helping bring about implementation agreements is the Network Management Special Interest Group (NMSIG) within the National Institute of Standards and Technology/OSI Implementor's Workshop. The NMSIG plans by the end of the year to have implementation agreements for 11 systems-management standards and a number of managed-object standards.

Another group working toward the same goal, but with more emphasis on telephone management, is the OSI Network Management/Forum. Even after each group has its say and vendors start making real products, there will have to be tests to ensure that one vendor's implementation works with that of another.

MANAGING THE FUTURE

Providing a comprehensive set of management tools is often hampered by the process of standard making itself. Because of the various interests involved, progress is likely to be slow. At an ISO plenary session held in Sydney, Australia, early in December 1988, ISO members spent hours on the French delegation's objection to the word "domain" in a standards document. Once it had been decided to keep the word, the group's recommendation had to pass on to a higher level committee, itself a subcommittee of the ISO executive body.

Sometimes one version of a standard may be inconsistent with another. By endorsing an early version of an ISO management standard instead of one that had been further refined, the Manufacturing Automation Protocol/Technical and Office Protocols (MAP/TOP) Users Group—which has helped influence which standards get implemented as products—lent support to an early, immature version that became obsolete before many vendors had incorporated the standard in their products. Having standards, moreover, will do little to help with the vast amounts of equipment and software bought before standards are implemented.

While they provide essential tools for managing different vendors' systems, some of the trickiest management tasks will never be standardized. The network manager must still judge what must be done to address a given problem, and the newness of large private networks means that few people have mastered the needed skills—a hybrid of data communication expertise, systems analysis, and the technician's craft.

Even with a good manager, network complexity sometimes outstrips human ability to interpret information available at the control center. Many in the field believe that artificial intelligence techniques will help to automate the monitoring and control functions. A few vendors have products that claim to incorporate expert systems. Most can be used only with a single vendor's products—or even just a single product within a product line. For the immediate future, however, standards and new technology will help provide a partial answer to the question: What is happening on the network? □

Standardizing Network Management for TCP/IP Environments

This report will help you to:

- Become aware of the work being done to establish TCP/IP-based network management standards.
 - Plan a migration path from TCP/IP network management to ISO network management.
-
-

The broad demand for multivendor networking environments has fueled the growth and acceptance of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol family into mainstream commercial environments. Networking vendors have joined forces with original members of the Internet community to help refine and evolve the protocol suite to meet the challenges imposed by its new user population. As a result of the explosive growth in size and complexity of TCP/IP-based local area network (LAN) and wide area network (WAN) environments, the topic of network management (NM) standards has received increased focus and attention. At the occasion of the TCP/IP Interoperability Conference held in March 1987 in Monterey, CA, a working group was formed under the auspices of the Internet Engineering Task Force (IETF) comprised of vendors, users and researchers with the following objectives:

- To define a framework for TCP/IP network management;
- To define a standard set of mechanisms (protocols) and management information in the form of Request For Comments (RFCs);

This Datapro report is based on "Standardizing Network Management for TCP/IP Environments," by Amatzia Ben-Artzi, Eric Benhamou, and Jim Robertson, Sytek, Inc. and Bridge Communications, Inc. from the IEEE 1988 Network Operations and Management Symposium, New Orleans, LA, February 28- March 2, 1988. © 1988 IEEE Communications Society.

- To provide a clear migration path from TCP/IP network management to ISO network management; and
- To conduct the specification work on an aggressive timetable with a completion milestone for draft RFCs of June 1988.

This report describes the work performed by the Network Management Working Group (NMWG) of the IETF, including the architectural model for network management systems, the management information structure, and the current status of the various specifications emanating from these activities.

ARCHITECTURE FOR NETWORK MANAGEMENT SYSTEMS (NMS)

Network management systems (NMS) contain the network management application level software and appropriate transport software. They can perform any number (zero to all) of these functions. They are independent systems, and should be able to interact with other NMSs, accept commands, and send information. In the context of network management, these systems are agents and/or managers. Agents are the systems that collect information and send reports to manager systems. Managers are associated with a human operator who initiates commands and receives reports. The manager station relies on different agents, but may be directly involved in the collection of data requested by the human operator or manager.

Standardizing Network Management for TCP/IP Environments

Although the ISO model does not fully specify the management relationships between multiple "Open Systems" with a global system, the NMS defines a hierarchy that uses a practical approach on a system construction. The hierarchy defined here allows for flexible relationships between agents and managers, and allows the system to grow naturally.

Agents constitute the lowest hierarchical level of network management. They can be present in every IP-addressable device of the network—terminals, terminal servers, PCs, hosts, bridges, routers, or gateways. However, the managed objects in each such IP device are not necessarily constrained to objects residing inside the device itself; they can be any object that has an Object-ID assigned to it, provided the IP-addressable device has agreed to be the object's "representative."

The agents report to, and are controlled by, a manager station. The association between an agent and a manager is not fixed, and can be configured. The number of managers, and the number of agents per manager, depend upon network configurations, the power of the manager stations, and other considerations of the human manager. If multiple stations have been defined, they can be bundled in a tree structure, so that a single station is configured at the top. This station is expected to be the central location where management of the entire network occurs. The number of manager stations reporting to the next level of the network management hierarchy is also not fixed, and is subject to considerations of performance and network configuration.

NETWORK SERVICES

Among the most expensive investments in developing network management solutions are the applications themselves. These are the ultimate processes by which the human operator is exposed to different views of the network. Applications can range from a single query/response process to a program capable of producing a full-color topology network map, with implied color indicating load levels of the different segments.

To protect this investment, a clear service definition is required. This definition must include the services an application can request from the underlying network, and the services provided by the underlying mechanisms (either as response to service requests, or on a voluntary basis).

Services provided by the NMS are grouped into four major areas.

- Association
- Monitoring
- Control
- Event Reporting

Each of these can be split into more detailed services.

- Association includes Initialize, Terminate, and Abort of Association
- Monitoring includes Get Operation
- Control includes Set, Confirmed Set, Action, Confirmed Action, Create and Delete Operations
- Event Reporting includes Confirmed and Unconfirmed Events, Structure of Management Information (SMI)

Another critical issue in the context of managing information is the ability to get to the right information and interpret the data in the correct way. In order to achieve this goal, two mechanisms are defined:

- Each "managed object" has an associated description. The entire collection of Object Descriptors is organized in a tree structure, so that traversing the tree in a systematic way leads to any desired object. At each level of the tree search, an "instance" can be provided. The instance is the means by which a descriptor is bound to a real object (for example, if an object exists in many copies, the instance will point to the copy of interest).
- "Managed information" is organized in a well-defined structure, so that when reference to the information is made (e.g., Get or Set), using the object's structure will fully define the data of interest. Such categories currently defined include Counters, Tide-marks, Gauges, Thresholds, and Log events.

MANAGEMENT PROTOCOLS

Management protocols are mechanisms by which information is moved within an NMS in a way that is understood to the nodes engaged in a "conversation." Such conversations can be carried over connection-oriented or connection-less transport, reliable or unreliable delivery services—thereby providing different levels of quality of service (QOS). The real importance is that both the agent and the manager will be able to "understand" each other and know what service to expect.

Standardizing Network Management for TCP/IP Environments

In the TCP/IP arena, the underlying transport can be either TCP, UDP, or Remote Transfer Protocol (RTP). The management protocol itself is the current Common Management Information Protocol (CMIP) ISO protocol carried over Remote Operations Services (ROS). Presentation encoding is ASN.1.

DISTRIBUTION OF SOFTWARE

Two different extremes need to be considered when distribution of software is evaluated in the LAN environment. In order to load many (maybe thousands of) nodes simultaneously, a multicast approach must be taken. In order to load software from a remote distribution center, a point-to-point download must be considered, as broadcast is inappropriate in an Internet environment.

The solution developed by the Network Management Working Group is an integral combination of the two different approaches that allows the user to adjust the services provided by the download services to specific environments (e.g., large broadcasting segments, small LAN segments, direct connections to WANs).

STATUS OF STANDARDS

The Network Management Working Group uses the familiar Internet Request For Comment (RFC) procedure as a vehicle to publish their standards activity. Several RFCs have been identified. Each will undergo at least two iterations of draft/review before the scheduled final stable version in the June 1988 time frame. After a reasonable implementation period, the RFCs will be reviewed for possible clarifying annotation. As of this writing, titles and RFC numbers have not been assigned, but a summary of each planned RFC is given below:

- The overview of network management of RFC AAAA will provide a summary of the purpose and goals of NM, the architecture/model, the management services provided by NM, and the protocols.
- The management services of RFC BBBB describe a generic set of service primitives, which are applicable to manage either Internet or ISO networks. These services provide for manager-agent association, monitoring, control, event reporting, access control, and synchronization.
- The management protocols of RFC PPPP describe the manager-agent protocols used to convey network management information. It consists of specifying how the ISO protocols CMIP and ROS operate in the TCP/IP environment.

- The software download protocol of RFC DDDD describes how download is accomplished in different configurations. It provides for multiple simultaneous loads and loading across gateways.
- The management information overview of RFC CCCC describes the structure of management information and the naming or identification of managed objects.
- Several individual RFCs are planned for specifying managed objects. Included will be RFCs for layer protocols (e.g., TCP, IP, ICMP, ARP, UDP, 802.X, X.25, etc.), and overall system parameters.

IMPACT AND FUTURE

The goal of the NMWG standard activity is to provide the fundamental mechanisms that enable multi-vendor, interoperable monitoring and control of the network. The impact of these standards will be network management applications from different vendors that have a unified view of the management information base and the capabilities to monitor and control all attached devices in their administrative domain. The NM applications, which constitute the bulk of the NM effort, will use and build upon the fundamental mechanisms provided in the standards. The NM applications will be vast and varied, with different user interfaces, distributed or centralized, and incorporating manual, automatic and/or artificial intelligent control. They will be designed to provide for the NM requirements of different users. Maintenance data and facilities will be provided to technicians for problem detection and diagnosis, installation and check-out, and preventative maintenance. Operations data and facilities will be provided to network operators for performance monitoring, configuration management, and network access management. Planning data and facilities will be provided to network designers for modeling, simulation, and network design.

In general, NM applications can be applied by having elements of the following categories of NM:

- *Configuration Management*: determination and control of the state of the network, both the logical and physical configuration of the network. Element examples are: port parameters, protocol timer values, physical location of nodes, bootload, threshold values, etc.
- *Performance Management*: control and assessment of the performance of the network. Element examples are throughput, error counts, retry counts, and historical statistics.

Standardizing Network Management for TCP/IP Environments

- *Fault Management*: the detection, isolation and correction of abnormal operation in the network. Element examples are audit trail, threshold exceeded events, protocol violations, layer confidence tests, diagnostic tests, and link down event.
- *Accounting Management*: the collection of data related to the charges or cost of use of resources in the network. Element examples are cost per packet and connection time cost.
- *Security Management*: the protection of and control of access to resources in the network. Element examples are authentication, encryption, access control, security logs, etc.

In looking to the future of NM in the Internet environment, another goal of the NMWG is to provide a migration path to ISO NM. Initially there will be exclusively TCP/IP nodes, then a long period of time with a mixture of both TCP/IP and ISO nodes, then ultimately, ISO homogeneous networks. The protocols are specifically different but have analogous functionality. The goal is to protect the large investment in NM application software. This is accomplished by the common management services specification to which the NM applications interface. These services are carried via CMIP/ROS protocols to manage ISO and TCP/IP nodes. Also, the structure of management information is common between the two and, where possible, the identification of parameters in analogous protocols are kept aligned. □

Managing TCP/IP-Based Networks: The HEMS Model

This report will help you to:

- Design and use internetwork management protocols.
 - Learn why standardizing on protocols is not enough to ensure interoperability.
 - Review the HEMS model and discover how it goes beyond protocols standards to support interoperability in heterogeneous environments.
-
-

HEMS¹ is an internetwork management protocol designed to work with the TCP-IP protocol suite.² While network management protocols are typically designed to manage a particular network using that network's link-level protocol to transfer requests, internetwork protocols are designed to manage a collection of interconnected networks that use different link-level protocols, and must rely on higher-level protocols (network-level or above) for connectivity.

In late 1986, a number of researchers, users, and vendors decided that the TCP-IP protocol suite needed a standard internetwork management protocol and that none of the then-existing management protocols was a suitable candidate for standardization. They formed working groups to develop new management systems and protocols that could be considered for standardization. HEMS, the High-Level Entity Management System, is the best known of the systems to come out of this effort. It expects each node on a network to support a virtual management database and provide database query language primitives that allow remote users to modify the database. In this report, the architects of HEMS present a detailed overview of the system.

This Datapro report is based on "The High-Level Entity Management System (HEMS)," by Craig Partridge and Glenn Trewitt. © 1988 IEEE. Reprinted, with permission, from *IEEE Network*, Volume 2, Number 2, March 1988.

MOTIVATION AND REQUIREMENTS

HEMS is designed to address the problems raised by several trends in IP networking. First, an increasing number of vendors are offering TCP-IP products and, as a result, IP networks are becoming more heterogeneous. A single Ethernet now commonly hosts TCP-IP nodes manufactured by half a dozen vendors, and users would like to be able to manage all of these nodes with a single management package. Beyond pointing out the need for standardization, this heterogeneity suggests that the standardization needs to be fine-grained. Standardizing on protocols is not enough; both the format and the meaning of management information exchanged must also be standardized. A network manager should not be expected to keep track of the vagaries of dozens of different implementations. HEMS addresses the problem of heterogeneity by requiring that management information have the same meaning on all nodes, and all information be exchanged using a common external data format.

Another trend is the explosive growth in the size and topological complexity of IP networks. The largest IP network, the Internet, now includes over three hundred networks and well over ten thousand hosts and is projected to grow by over forty percent in the next year. (This is a deceptively low count of networks, since the IP architecture permits several physical networks to be combined into a single logical network using a procedure called *subnetting*.)

Managing TCP/IP-Based Networks: The HEMS Model

Increased size has several consequences. Because network paths are, on average, longer and more complex, data loss due to problems such as congestion, link failure, and router error becomes more likely. (In the TCP-IP protocol suite, the network layer is not reliable.) The increased likelihood of loss suggests that management systems should use a reliable transport protocol to ensure that management information is reliably exchanged. Another reason for using a reliable transport protocol is that when the network is broken, we want to use the management protocol to fix it, and reliable transport protocols are usually more robust and better able to ensure that management requests reach their destination.

The trend toward greatly increased network size also means that IP networks are becoming too large to be managed by a single management center. On the Internet, we can see movement toward shared network management in which several management centers collectively manage an internetwork, where each management center takes responsibility for a section of the internetwork. The sections frequently overlap; for example, the centers at two organizations may share the responsibility for a router that connects their networks. As a result, a management protocol must allow a node to be managed by more than one manager. The Internet also has a tradition of allowing researchers to monitor network devices, which implies the need at least for multiple monitoring applications to be able to access a node concurrently.

Even with multiple management centers, most centers will be managing dozens of networks and hundreds or thousands of nodes. These centers probably cannot be expected to actively search for network failures by polling or maintaining continuous connections; some mechanism must be provided for nodes to notify their management center (or centers) that a problem exists. HEMS allows nodes to send unsolicited reports called *events* to management centers.

One final design constraint was a strong desire to minimize the cost of implementing the management system on network nodes and, in particular, to minimize the cost for routers and other dedicated network equipment. By cost, we mean both implementation effort and the overhead imposed on a network node (that is, code size and processing demand). The desire to minimize implementation effort was pragmatic; vendors are more likely to support a system that is inexpensive to implement. But the desire to limit processing costs was more technically motivated. Router technology is having some difficulty keeping pace with increasingly available high-speed communications links. A management system that continuously consumes even a small percentage of a router's processing capacity may have a noticeable impact on

the router's ability to keep pace with offered traffic load. HEMS tries to reduce the per-system cost by making the code that processes management requests very simple, supporting a small number of primitive operations. Consequently, the application that requests management services must provide the necessary intelligence to use primitive operations to support higher-level management operations.

OVERVIEW

HEMS is designed to operate in the TCP-IP suite of protocols. Since IP is the first level of uniformity in the TCP-IP architecture, each managed node must have an IP address. Systems without an IP address, such as link-level bridges, may be managed through proxy or translator nodes that do have an IP address. (The difference between a proxy and a translator is that a proxy node is a single designated manager of a particular device, whereas more than one translator could be used to relay a request to a single link-level device.) Applications see the managed data as a hierarchical database and use primitive operations to traverse, read, and modify it. The values in the database correspond to state information in the node.

To monitor the node, applications use read operations, which extract data from the node, package it in external format, and return it to the application. To control the node, applications use operations to write, create, or delete values in the database. These requests are translated into appropriate control operations on the node. For example, deleting a routing entry in the database causes the corresponding entry in the nodes' routing table to be deleted. The software that manages the database and processes the operations sent by applications is called the *agent*.

Applications and agents communicate using *messages*. To request management services, an application sends the agent a query message containing one or more database operations. The agent responds to the query with a reply message containing the monitoring information requested and the completion status of any control operations. Messages are transmitted using the High-Level Entity Management Protocol (HEMP).

Events are treated as unsolicited HEMP reply messages. Applications that wish to receive events deposit their addresses on a per-event basis in the database. When an event is generated, it is sent to all the applications that are on its list.

To ensure that information is reliably exchanged, HEMP uses a reliable transport protocol for transferring messages. Because HEMP is message-oriented,

Managing TCP/IP-Based Networks: The HEMS Model

the most suitable transport protocol would be one designed for transaction processing (for example, the Versatile Message Transport Protocol, VMTP).³ Unfortunately, no transaction protocol is yet considered a standard TCP/IP transport protocol, so HEMS currently uses a mix of transport protocols: the Transmission Control Protocol (TCP) for query-reply exchanges, and the User Datagram Protocol (UDP) for events.

EXTERNAL DATA REPRESENTATION

The data on which HEMS operates is drawn from a rich set of primitive types and may have arbitrarily complete structure. Because of this, the external representation of the data must be very flexible, providing tags to identify the type of the data. By attaching a tag to each data item, some information about the meaning of the data is passed along with the data. In addition, the length of data items must be explicit. Giving the length of data items allows a program to handle the data without understanding anything about it. Therefore, each piece of data is conceptually a tuple of `<type,length,data>`.

Tags are actually small integers, even though we write the names associated with them. The meaning of some tags, for example, well-known types such as Integer and Octet String, is universally understood. Most tags, however, must be interpreted in context and have different meanings depending on where they are found.

HEMS requires that the external representation handle both simple data, such as Integers and Octet Strings, as well as structured data, composed of one or more data items, each encoded in the representation. With such a representation, it is possible to encode any object that one might define in a language such as Pascal or C.

For the purposes of this report, we will ignore the length of data items, representing a simple data item with the notation `tag(value)` and structured data with `tag{item1, item2, ...}`, where each item may be either simple or structured.

The current definition of HEMS specifies that the data will be externally represented using the ISO Abstract Syntax Notation 1 (ASN.1).⁴ ASN.1 provides the facilities described above, although any comparable external data representation could also be used.

DATA ORGANIZATION

We met several design goals by carefully designing the data organization in the managed entities. In most systems, the data to be managed is already present, in operating system tables, the file system, hardware PROMs, or somewhere else. However, there is rarely a uniform way to access this data. The primary task is to give some structure to this data and provide a common access method. The organization should be extensible so that new system components, such as protocol modules or hardware interfaces, can be added later. Finally, the system for naming the data in the entity should be straightforward and compact.

We choose a hierarchical organization for the data. Each logical collection of data, such as information about a hardware interface, a routing table, or a protocol module, is grouped together at one node in the tree. The root of the tree defined by HEMS is "Information About an IP Entity." The descendants of the root node correspond to different layers of the IP suite and different applications within a layer.

Data in the tree is named by giving the path from the root to the desired node; the resulting name is similar to a pathname in a hierarchical file system. A name, in the external representation, is merely a piece of data with no value. We represent this by just giving the tag name, with no parentheses or braces following. Therefore, the name of the entity's clock, which is contained in the *SystemVariables* node would be *SystemVariables{clock}*. Since many values can be represented in one data item, many pieces of data in the tree may be named in one "name." For this general case, we use the term "template." The template *SystemVariables{clock,name}* names two items in the same node. With deeper nesting, templates can span several nodes, although we expect such usage to be infrequent.

Much of the managed data is found in tables, such as routing tables, lists of interfaces, and address mapping tables. Only a few of these tables have useful indices; most are unordered sets of tuples. Since these tables can be large, it is desirable to be able to select some subset of the table entries for an operation. This is done using *filters*, which make a simple comparison of data in each table entry to decide if it is a candidate for an operation. Filters are encoded directly in ASN.1 as expression trees. This representation is compact and easy to evaluate.

Overview of Data Hierarchy

Figure 1 shows the upper nodes of the data hierarchy. The descendants of the root node correspond to logi-

Managing TCP/IP-Based Networks: The HEMS Model

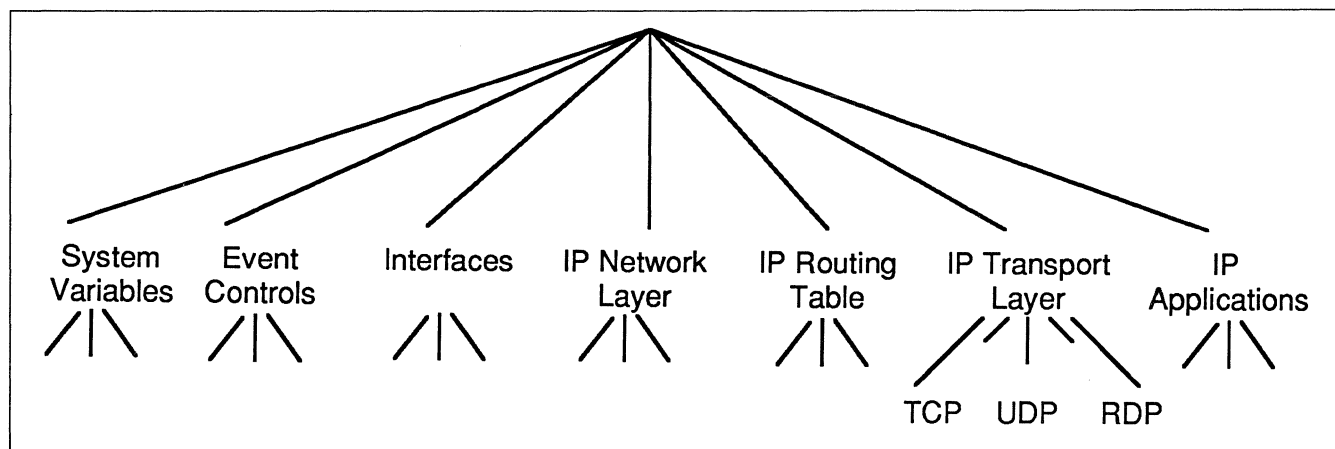


Figure 1. Upper nodes of the data hierarchy.

cal divisions of the functionality of the entity. Most of these nodes correspond to particular levels of the seven-layer model. For example, the tree below the IP TransportLayer node contains information about transport layer protocols such as TCP, the Reliable Data Protocol (RDP), and UDP. The one exception to the one-node-per-layer rule is the routing section. While it is technically a network-layer function, routing tends to cross layer boundaries; for example, routing information is commonly exchanged using a transport protocol. Some layers, such as the presentation layer, are not represented because they have nothing to manage as yet.

A few other nodes represent classes of management information about the entity which do not correspond with any of the seven layers. The SystemVariables node models information about the system software that runs on the node. This information can be critical for network management (for example, the control values to reboot the node are stored here) but do not fit into the seven-layer architecture. Another exception is the EventControls node, below which is stored all event management information.

The HEMS protocols will work with other data hierarchies; this is just the data hierarchy designed for use on nodes in TCP-IP Internetworks. Furthermore, in some cases the information stored below some nodes is not yet defined. Although a place exists in the hierarchy for information about applications, for example, no management information for applications has been defined. While we believe it is useful to be able to manage applications, so far no one has suggested an application that needs remote managing, much less an abstract representation of an application that might appropriately be placed in the hierarchy.

Extensibility

An important goal of the HEMS design was to make it possible to extend the data tree without requiring extensive reimplementations of the HEMS systems. There are two different aspects of extensibility. When new protocols are defined, such as bulk data transfer protocols, a new subtree will be added to the definition of the tree. Anyone who implements the protocol will be able to make it manageable by adding appropriate code to the agent. Similarly, someone who wants to manage the new protocol will be able to do so by extending their management application. This clearly requires that someone write management specifications for new protocols, which hopefully would be an integral part of a protocol definition.

A more difficult situation arises when a vendor wants to manage something peculiar to its implementation. This may be just one or two pieces of data associated with an already-existing node in the tree. HEMS provides an escape mechanism so that an implementation may add arbitrary additional data to any node in the tree. If necessary, entire subtrees can be added, although standardizing a definition would be preferable. At any node in the tree, a subnode named "VendorSpecific" may be added. VendorSpecific is a well-known name—it will be recognized anywhere in the tree, and there may be several at different places in the tree (but not more than one at a single node). A VendorSpecific node can contain any information at all. By giving such nodes a well-known name, application software that retrieves data from a VendorSpecific node can recognize what it is and use other HEMS mechanisms to find out more about the data.

Data Attributes

It is possible to understand the meaning of the information in the data tree by knowing the correspon-

Managing TCP/IP-Based Networks: The HEMS Model

dence between the tags for the data and the definition of the data. Typically, this will be standardized in a monitoring specification. If every monitoring center has complete, up-to-date specifications, this works well. Often, however, the specifications will be unavailable or out-of-date, or the monitoring center may be dealing with an out-of-date entity.

For each node in the data tree, HEMS defines a collection of *attributes* that describe the node. An *AttributesInfo* structure contains descriptive information about the node, such as the base data type (Integer, Counter, Octet String, and so forth), a short description of the data's meaning, a very short string to be used as a label, the units of the data, and miscellaneous properties of the data. Most of this information will be shared among data items of the same type, leaving only the descriptive information to be filled in on a case-by-case basis.

By making attribute information available directly from the monitored entity, a monitoring application with little or no specific information about the data tree can let a human operator browse through the data tree, producing meaningful displays of the information. This capability may be crucial when Vendor-Specific data, for which no other documentation is available, is being examined.

Low Level Data Types

Defining the overall structure of the database was only half the problem. We also had to determine what data types were required to properly represent the information being managed. ASN.1 defines several basic types, such as integers and sequences of octets, and also provides mechanisms for grouping basic types into aggregate structures such as sets and sequences. But some additional types were also required.

The additional data types generally fell into two categories. Some types were truly new; they represented basic abstractions not defined by an ASN.1 basic type, for example, the *IpAddress* type, which represents an IP address, and the *Counter* type, which represents a roll-over counter. The other data types were aggregates built from basic types. Examples of such aggregates are a type to represent histograms and the *VendorSpecific* aggregate.

Choosing proper representations proved to be difficult. The development of the Counter type is a good example of the problems involved. We knew that counts of various functions performed (for example, total packets switched or routing updates received) are an important network management tool. Further-

more, experience shows that on machines with common integer sizes, such counters often overflow. As a result, we needed some abstraction that incorporated the notion that a count had reached a limit; a simple integer type would not do. In the past, some systems have used latch counters, that is, counters that freeze at their maximum values instead of rolling over to zero. Unfortunately, latch counters must be reset. This is at odds with supporting multiple management applications. Who gets to reset the latch counter? How do we ensure that the latch counter always gets reset so all applications can continue to read it? We therefore chose to use roll-over counters, which cannot be reset. Roll-over counters are always readable, do not require maintenance, and if they are made large enough, rollover only infrequently (once every few days).

QUERY LANGUAGE

Conceptually, a HEMS query consists of a sequence of tagged data items designed to be interpreted by a simple stack-based interpreter. (Although the query language is defined in terms of the operations of a stack machine, it does not have to be implemented with one.) One tag is reserved for operation codes; all other tags are for data that will eventually be used by an operation. Nonopcode data items are pushed onto the stack when received. Opcodes are immediately executed and may remove or add items to the stack. Because the external representation is tagged, very little needs to be done to make the incoming data suitable for use by the interpreter.

There are eight operators in the query language. The overall operation of most is very similar, with differences only in the final effect (retrieving vs. setting data values). All of the operators work in the context of a current location in the data tree.

The BEGIN operator takes a path from the current node in the tree to a lower node, and makes that the current node. The END operator undoes the effect of the most recent BEGIN.

The GET operator takes a single template and retrieves the data named in the template. It "fills in the blanks" in the template. The GET-ATTRIBUTES operator retrieves the attributes for the data named in the template.

The SET operator is identical in form to the GET operator except that, rather than taking a template as an argument, it takes a template with values already filled in and sets the named data in the tree to the given values. The CREATE operator takes a filled-in template and creates a new entry in an existing table.

Managing TCP/IP-Based Networks: The HEMS Model

The BEGIN, GET, SET, and GET-ATTRIBUTES operators can be given an additional *filter* argument to restrict the operation to selected entries in a table. A filter is a Boolean expression that is executed for each entry in a table, making comparisons on the data in the entry. The operations are restricted to comparisons of constants and AND, OR, and NOT expressions. If the filter is true, the operation is performed on the entry, otherwise the entry is ignored. One additional operator, DELETE, removes entries that match a filter from a table. Unlike the other filtered operations, BEGIN uses the first table entry that the filter accepts, rather than iterating over all of them.

Control Operations

Control operations are performed by setting data items to values that trigger special operations. For example, the SystemState data item might receive values that halt, reboot, or reinitialize the operating system. This virtual "command-and-status register" model allows control operations to be extended in the same way as the data tree.

Monitoring Information

A query is usually used to retrieve information from the data tree on the remote entity. The result of a query is simply the data requested, returned as a composite object in the external data representation. That is, the result is a selective walk of the data tree. The overall skeleton of the tree is determined by the sequence of BEGIN and END operations; GET operations are used to fill in the nodes of interest.

This scheme is extended to the other operations. SET and CREATE return the value that was actually put in the tree. DELETE returns an empty node if the data was actually deleted, otherwise it returns the data that could not be deleted.

Two Query Language Examples

To illustrate how the HEMS query language can be used, we present two examples. Both examples involve routing.

In the first example, suppose that we discover that the routing system on our local gateway has been misconfigured to use the wrong routing metric to interpret routing information. To fix this problem we would like to change the type of metric used (from, say, type 0 to type 1). At the same time we would like to get a

complete listing of our routing table to discover if our current routing information has become corrupted as a result of this error.

In HEMS, the routing information is stored in the IpRoutingTable dictionary. The metric used is in the metricUsed field, while the set of all routing entries is stored in the RoutingEntries field. The query to read the routing table and set the metricUsed field is shown below in its symbolic representation. (It would actually be sent as the twenty-byte ASN.1 encoded hexadecimal sequence:
5f2004101038301014101028400410101410104)

```
IpRoutingTable BEGIN
metricUsed(1) SET
RoutingEntries GET
END
```

The agent would reply to this query with a message containing the following data:

```
IpRoutingTable{metricUsed(0),
RoutingEntries{... all the routing entries ...}}
```

For the second example, suppose that we are trying to locate a routing problem which appears to affect our connectivity to two networks: 128.89.0 and 192.5.58.0. To try to locate the problem, we want to send a query to our local gateway, asking it what routes it currently uses to reach these networks.

Because we are interested in only two routing entries we should use filtered operations to extract the particular routing entries we want. The filter used here is an *or*-filter, which selects all instances of the RoutingEntry structure which have a destination network of either 128.89.0 or 192.5.58.0. The query looks like this:

```
IpRoutingTable{RoutingEntries} BEGIN
RoutingEntry
Filter{ or { item { equality { routeDst(128.89.0) } } }
      { item { equality { routeDst(192.5.58.0) } } } }
GET
```

If there was one route for each network, the answer might look like this:

```
IpRoutingTable{ RoutingEntries { RoutingEntry
      { routeDst(128.89.0), ... },
RoutingEntry { routeDst(192.5.58.0), ... } } }
```

EVENTS

Nodes can spontaneously report potential problems to one or more management centers by using *events*.

Managing TCP/IP-Based Networks: The HEMS Model

While a network could be managed exclusively by polling critical nodes on a regular basis, experience suggests that this is inefficient. Most of the polls are unnecessary (because the node is running properly) and it often takes some time to get around to polling a node after it develops a problem. Allowing a node to report a problem directly is more effective.

When a node detects an anomalous condition (for example, a routing loop or a defective link), it triggers an internal event and sends an event message to any management centers that have requested that they be notified of the particular event.

Each event is assigned a two-part value. The *event code* is a number that indicates the general type of problem. Event codes have the same meaning across all nodes; for example, event code 100 might mean "routing loop detected." The *event index* is an implementation-specific value that vendors may use to further classify the event. Event code 100 with an event index of 27 might indicate the routing loop occurred using a particular routing protocol. The event message also contains a text description of the event that provides more context information to the management centers that receive the event.

Finally, each event message contains optional data from the management database. This data is programmable; management centers load small queries into per-event query buffers and when the event occurs this query is run.

Any data it extracts is included in the event message. Centers can also specify the frequency with which they receive an event: every time or once every *n* events that occur. Management centers can thus customize their event messages to allow them to better isolate problems.

One problem with the event mechanism is the possibility of "event floods," a condition in which a malfunction causes a system to send a continuous stream of event messages. HEMS handles this problem by keeping track of how frequently an event message has been sent and suppressing the messages if the transmission rate exceeds a threshold.

MESSAGES AND HEMP

An important characteristic of HEMS is that management interactions are exchanges of messages, with a message typically containing several operations or several pieces of monitoring data. This interaction contrasts with a typical remote procedure call (RPC) mechanism, in which each operation or data item is sent separately.

We choose the message-oriented architecture for performance reasons. We believe that because HEMS query language operations are fairly primitive, most management interactions will involve performing more than one operation. If more than one operation must be performed, it is substantially more efficient to send the operations in a group rather than individually.

Grouping operations allows requests to be effected faster. If each operation depends on the success of the previous operation (which is often the case in HEMS), then an RPC mechanism is very slow because before it can send an operation, the requesting application must confirm that the previous operation has been received, processed, and acknowledged by the remote agent. If the network delay is long, then the application may spend much of its time waiting for acknowledgments. In contrast, a message system sends all the operations together and is less affected by network delays.

Grouping operations also allows all state information such as access control, request identifiers, and encryption information to be transferred and processed once for each operation group instead of once for each RPC operation.

The HEMS message protocol is called HEMP. Each HEMP Message contains zero or more optional requests for special services followed by a common header, which contains information about the type of message, the version of the protocol in use, and a message ID that allows requests and their corresponding replies to be matched. Following the header is the message body, which contains the management information (i.e., the actual query or data). There is no limit to the length of a HEMP message.

HEMP supports four types of optional special services: authentication, access control, encryption, and accounting.

It is often necessary to confirm that a management request is from an authorized user, or that the management information comes from its claimed source. Before an agent shuts down its node, the agent should be able to confirm that the shutdown request came from an authorized network management center. Similarly, if a management center receives an event indicating that a key node has failed, it might wish to confirm that the event came from the node and not from another, possibly malicious, source. HEMP provides access control and authentication services to meet these needs.

Encryption services must be available because some management information is sensitive. A good exam-

Managing TCP/IP-Based Networks: The HEMS Model

ple of sensitive information is traffic information: many organizations (for example, defense agencies) may not want outsiders to learn about traffic patterns within their networks. While access control ensures that only authorized users are entitled to ask for information, it does not protect against eavesdroppers. That protection is provided by encryption.

HEMP also provides accounting services because a few users have suggested that they may want to do per-message accounting. We suspect, however, that accounting is better handled at another level.

COMPARISON WITH ISO

The International Standards Organization (ISO) is currently attempting to develop standards for network management. Much of the effort so far has been aimed at defining an architecture for management, parts of which are very similar to HEMS, and other parts of which are quite different. In most areas, HEMS can be integrated smoothly with the ISO architecture.

At the lowest level, both the ISO architecture and HEMS use the same data-encoding standard for exchanging information: ASN.1. The basic data types provided by the ISO architecture and HEMS are essentially identical; differences exist largely because the ISO definitions are still evolving. We have tried to keep HEMS close to the emerging ISO data-type definitions.

The ISO committees have structured the management data into a hierarchy referred to as the "Management Information Base" (MIB). It is similar in philosophy to the HEMS data tree, although only the structure (a hierarchy) is defined in the existing draft standards. The way that data is named, which we feel must be standardized, is left to be defined in the specification for each protocol layer (corresponding to a subtree in the MIB). For example, one proposal specifies that each item is named with a unique integer, while another suggests naming with text strings, like file names. Both of these may be adopted because each can apply to different subtrees of the MIB. The HEMS data tree fits into the MIB in the niche reserved for IP Information.

The ISO and HEMS architectures differ significantly in the way the services they provide are implemented. ISO uses a "Common Management Information Protocol" (CMIP) which implements a "Common Management Information Service" (CMIS).⁵

CMIP is built upon an RPC layer known as "Remote Operations Services" (ROS). As stated earlier, we

chose a message-based, query-response model for management rather than RPC. While this is a fundamental difference that prevents a one-to-one mapping between the two systems, the services provided are very similar, and it should be possible to build translators between the two.

HEMS does provide some features that CMIS does not. HEMS explicitly defines how the management data is structured and named, and how it may be extended; CMIS leaves most of these decisions to the designers of each subtree or protocol layer. We believe that defining a consistent infrastructure for the management information will make it easier to define additions to the database. In addition, CMIS has no notion of data attributes; only the data itself can be retrieved. In HEMS, it is possible to discover additional information in the database beyond what is defined in standards. It is also possible to discover the meaning of this data.

The definition of HEMS filters was derived from a recent draft of the CMIP specification. It now seems likely that the CMIP definition of filters will be changed before the specification is standardized.

CONCLUSION

HEMS has been implemented.⁶ Tests with an experimental HEMS agent have verified that the architecture of the system is sound and that the database architecture allows us to build a compact and inexpensive implementation, while achieving a powerful management interface.

ACKNOWLEDGMENT

We would like to acknowledge the considerable assistance of Charlie Lynn, whose comments on the earliest designs of HEMS were extremely helpful.

REFERENCES

- ¹C. Partridge and G. Trewitt, "The high-level entity management system (HEMS); RFCs 1021-1024," *Network Working Group Request for Comments, no. 1021-1024*, Network Information Center (NIC), SRI International, Menlo Park, CA, Oct. 1987.
- ²J. Feinler, O.J. Jacobsen, and M. Stahl, *DDN Protocol Handbook, Volume Two, DARPA Internet Protocols*, DDN Network Information Center, SRI International, Menlo Park, CA, Dec. 1985.
- ³D.R. Cheriton, "VMTP: A transport protocol for next generation communication systems," *Proc. of SIGCOMM '86*, Association for Computing Machinery, Aug. 1986.
- ⁴*Information processing systems—Open Systems Interconnection—Specification of Abstract Syntax Notation One (ASN.1)*, International Standard, No. 8824, International Organization for Standards, May 1987.
- ⁵*Information processing systems—Open Systems Interconnection—Management Service Definition*, International Organization for Standards, draft of Oct. 1986.
- ⁶C. Partridge, "A UNIX implementation of HEMS," *Proc. of the 1988 Winter USENIX Conference*, Feb. 1988. □

AT&T Unified Network Management Architecture (UNMA)

This report will help you to:

- Evaluate one of the first practical product implementations of OSI Management drafts and standards.
- Discover how you can use AT&T's Integrator in SNA networks.
- Compare AT&T's network management approach to IBM's.

Unified Network Management Architecture (UNMA) is AT&T's premier strategy for providing integrated, end-to-end management of voice and data networks in multivendor environments. The ACCUMASTER Integrator, announced on January 31, 1989, is the key component of AT&T's UNMA strategy. (See Figure 1.)

UNMA is a three-tiered architecture that follows the ISO/CCITT Open Systems Interconnection (OSI) standards and management framework, to the extent they are defined. (See Figure 2.) AT&T's approach of capitalizing on the increasing acceptance of OSI-like systems contrasts with the direction IBM took in developing NetView several years ago. IBM created its own de facto standard, designing NetView (and its predecessors, NCCF and NPDA) to manage networks conforming to SNA—a proprietary architecture developed primarily to connect IBM systems. Rather than create its own de facto standard from scratch, AT&T chose to use OSI Management drafts. OSI drafts are still subject to change, however, and leave a lot to be defined by the implementors. If successful, AT&T could create a de facto "implementation standard" of its own, as it were, using OSI Management drafts as the foundation.

Due to IBM's massive installed base of SNA networks, NetView's market presence will continue to dwarf OSI-type systems in the short term. SNA will never go away, and AT&T does not expect that it will. Rather, AT&T is aiming at the growing demand for multivendor network management products—in particular, for systems which are modeled on OSI drafts. As these systems evolve, AT&T hopes that innovative users will

support its implementation of OSI Management standards over Digital Equipment Corporation's and others and provide enough weight to sway the OSI market in the AT&T direction. To date, the momentum behind AT&T's implementation is outpacing Digital's EMA. To make matters more interesting, in September 1988, IBM announced that NetView will support OSI by mid-1990. While this support is convoluted at best, it again pits the two giants in an old battle, but in a new market—OSI Management.

Both NetView and UNMA provide a degree of integration by allowing the user to monitor both logical and physical aspects of the network from one central interface. (Network management systems providing a *logical network view* measure the network as represented by actual traffic passing over it. Systems controlling the *physical network* monitor and control the actual

Index to This Report	Page
The Structure of UNMA	103
Network Management Protocol (NMP)	104
Vendor Support for NMP	107
The OSI NM/Forum	107
Cincom's Role in UNMA	109
The ACCUMASTER Integrator	110
UNMA Functions	113

AT&T Unified Network Management Architecture (UNMA)

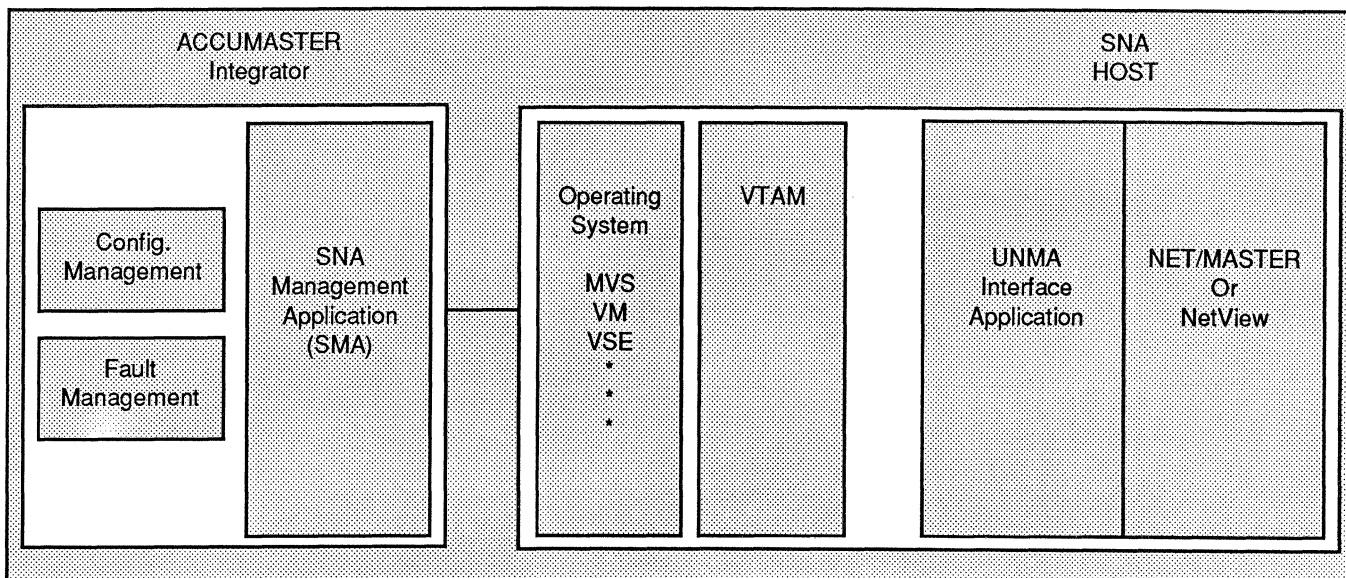


Figure 1. The Integrator provides an end-to-end view of both logical and physical elements in an SNA network. Cincom's UNMA Application works in conjunction with either NetView or Net/Master to extract SNA logical configuration and alarm information. The UNMA Application then employs Network Management Protocol (NMP) messages to pass this information to the Integrator's SNA Management Application.

circuits and network nodes.) Figure 3 illustrates the difference between logical and physical network management.

UNMA starts with the ACCUMASTER Integrator on the physical side and (in its initial implementation) uses 3279 emulation to provide cut-through capabilities to components on the logical side, such as IBM's NetView or Cincom's Net/Master. (At present, AT&T has no plans to offer its own logical management package for SNA networks. Cincom is, of course, hoping that ACCUMASTER customers will choose Net/Master over NetView.) In contrast, IBM's approach is to start with NetView on the logical side and integrate physical network management through its Link Problem Determination Aid (LPDA) and NetView/PC. Similarly, Digital's EMA starts with the Executive control program on the logical side and provides Access Modules that communicate to remote hardware using OSI-based protocols. (For more information on Digital's EMA, see Report NM40-325-101, "Digital Equipment Corporation Enterprise Management Architecture.")

UNMA provides for managing both data networks and voice networks. Its real edge, however, is its support for managing the customer-allocated portion of the public network. With UNMA, users can tie together three network management domains: the customer premises, the local exchange network, and the interexchange network. AT&T can provide UNMA customers with the means to integrate information from public networks with data from their private net-

work management systems. IBM's NetView and Digital's EMA stop short of the public network boundary and instead offer much more sophisticated logical network management capabilities, particularly at the applications level.

By building upon the evolving OSI standards and drawing from the company's unparalleled expertise in managing the world's largest network, AT&T is in an excellent position to offer an integrated solution for managing voice and data networks—in the future. The question is, when? OSI Management standards have not yet solidified; final approval and, more importantly, real OSI products are two years away, at best. AT&T seeks to offer an interim OSI-like solution right now and may convince users that OSI is far enough along to make an interim product useful. The major stumbling block to widespread acceptance of UNMA within the next 5 to 10 years is not that users must wait for OSI but rather, the *weight* behind IBM's massive installed base of SNA networks may overpower the influence of UNMA. SNA allows NetView to enjoy a visibility in the market that presents an enormous challenge to UNMA and the ACCUMASTER Integrator.

SNA predominates in the United States more so than in Europe, where government mandates are pushing private enterprises to embrace OSI. U.S. businesses competing on an increasingly global scale are beginning to feel the impact of the European OSI movement in a very real way. While this influence is carrying over into this country, the greatest initial demand for OSI

AT&T Unified Network Management Architecture (UNMA)

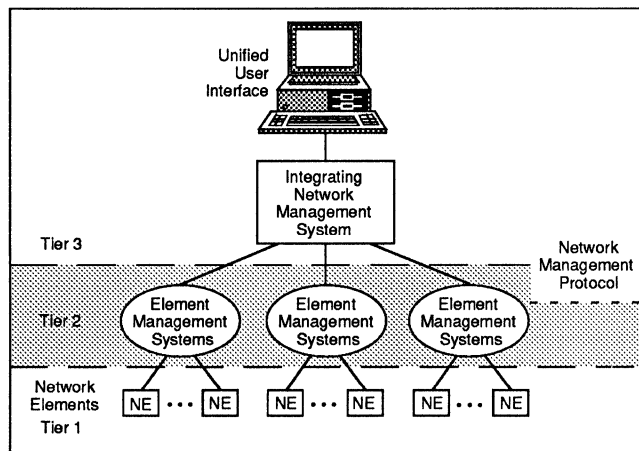


Figure 2. AT&T's UNMA features a three-tiered architecture. Tier 1 is composed of Network Elements (NEs), such as modems, muxes, PBXs, T1 lines, or computers. Tier 2 consists of Element Management Systems (EMSs), which may include modem- or PBX-based management systems, LAN management systems, or public network-based management systems. Tier 3 is the ACCUMASTER Integrator, announced on January 31, 1989.

products in the U.S. will come not from private companies, but from the U.S. government itself—as a result of the GOSIP mandate. In the long term, the GOSIP influence, combined with the gradually increasing user preference for standards-based systems over proprietary networks, may eventually foster a different climate in the network management market. This climate will be less conducive to the growth of strictly proprietary solutions in their present form (such as NetView) and provide a window of opportunity for both AT&T and Digital. Undoubtedly, IBM is gearing up to prepare for this future climate as well. Datapro believes that UNMA is designed to thrive in this climate and may do so if AT&T can effectively deliver its goods to market at the right time, with the right approach, and with a comprehensive physical/logical, public/private network management solution that IBM, so far, cannot offer.

The remainder of this report describes UNMA in detail, highlighting UNMA's Network Management Protocol (NMP) and ACCUMASTER product announcements. This report also discusses current vendor support for UNMA, including the OSI Network Management/Forum and AT&T's agreement with Cincom Systems.

THE STRUCTURE OF UNMA

UNMA features a three-tiered architecture with a Unified User Interface as its focus (see Figure 1). Machine-to-machine interaction occurs between the three tiers; human-to-machine interaction occurs at the Unified User Interface.

Tier 1 is composed of *network elements*, which may include customer premise equipment (CPE) such as modems, multiplexers, LANs, hosts, and PBXs. Local exchange carrier (LEC) networks, interexchange services, PTT, or international network services are also categorized as network elements.

Tier 2 is composed of the *Element Management Systems (EMSs)*, which manage network elements. An EMS provides what may be called "local" management capabilities—operations, administration, maintenance, and provisioning functions of managing a particular network element group. Today, most large networks include multiple EMSs, since vendors have traditionally supplied different systems for different products and services. It is not unusual to find different EMSs for computer hosts, matrix switches, T1 resource managers, and Ethernet LANS—all within the same corporate network. (See the section entitled "Currently Available AT&T Element Management Systems" for more information on AT&T EMSs.) Each Element Management System operates, all too often, as an island unto itself. The inability of disparate EMSs to share information prevents the user from obtaining an end-to-end view of the network.

Tier 3 of UNMA is designed to provide the desired end-to-end view by supporting communications between EMSs and the ACCUMASTER Integrator—the heart of the UNMA strategy. The ACCUMASTER Integrator communicates with individual Element Management Systems through a common protocol stack. This stack, called the Network Management Protocol (NMP), is AT&T's implementation of OSI management specifications as they exist today. (See the section entitled "Network Management Protocol" for more information.) Under UNMA, manager integrator systems also communicate among themselves via the standard protocol stack (Tier 2 to Tier 3). UNMA allows communications based on native (proprietary) protocols between the Network Elements and their respective Element Management Systems (Tier 1 to Tier 2). Communications based on standard protocols between Tiers 1 and 2 may evolve in the future, however.

The ACCUMASTER Integrator provides cut-through capabilities to the Element Management Systems, allowing users to exercise full capabilities of each EMS. This feature gives AT&T's product an edge over IBM's NetView. (See the section entitled "UNMA Products and Services" for more information about Integrator product features.)

Also located at Tier 3, the **Unified User Interface** creates the image of a virtual network, so to speak—providing a snapshot view of one large network (or subnetworks) that is actually composed of diverse entities. In the ACCUMASTER Integrator product, this

AT&T Unified Network Management Architecture (UNMA)

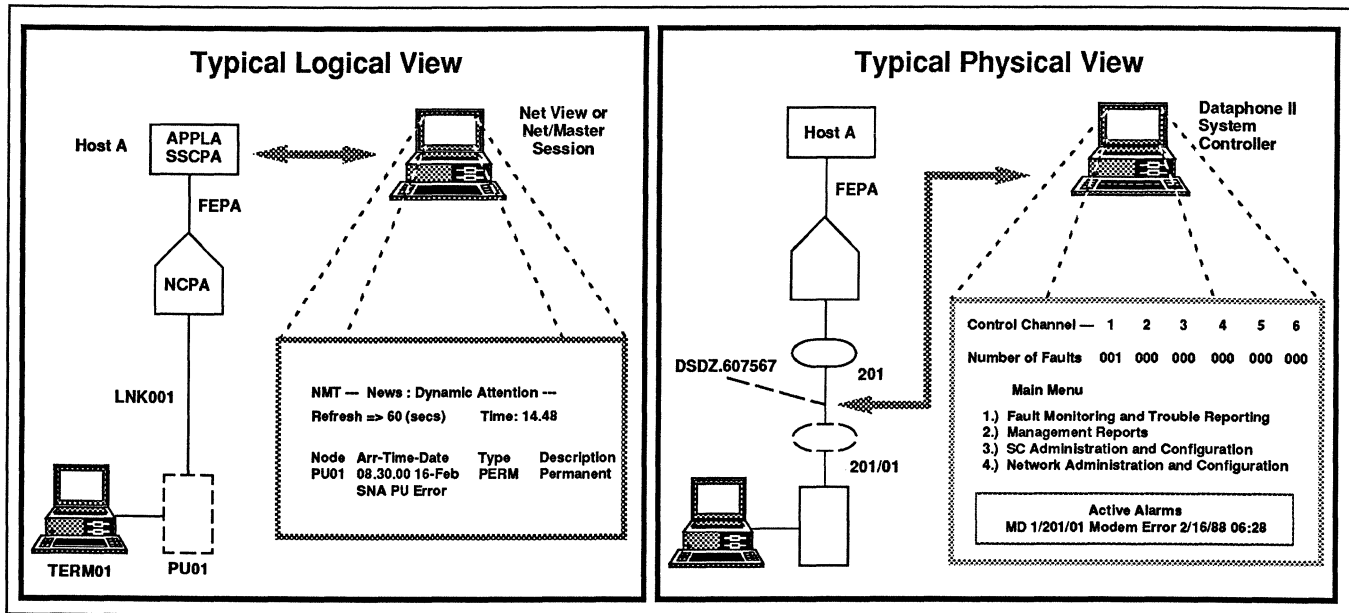


Figure 3. A typical logical view of network management (top box) sees the network in terms of the traffic passing over it. In SNA networks, management systems generate messages based on information from the host-resident System Services Control Point (SSCP) in ACF/VTAM. Device errors are traced to Physical Units (PUs), which SNA defines as a set of services performed by a node rather than an actual physical device. The typical logical view tends to trivialize the connection between the cluster controller and network devices. In a typical physical view of network management (bottom box), physical network management systems monitor the actual circuits and lines in the network. Alerts notify network administrators of device or line faults but do not provide information about applications affected.

is a graphics-based interface (a Sun Workstation) that presents integrated management data to the human operator.

NETWORK MANAGEMENT PROTOCOL (NMP)

As of June 1989, AT&T had published four NMP documents (TR54004 through TR5407) to help other vendors implement the interface between their proprietary Element Management System products and the ACCUMASTER Integrator. AT&T's work on NMP, as described in these documents, provides the most insight available to date on the practical mechanics of implementing OSI network management.

NMP is based on the OSI seven-layer reference model. (See *Datapro Reports on Communications Software*, Report CMS20-010-301, "ISO Reference Model for Open Systems Interconnection [OSI]"). NMP is also consistent with the OSI Management Framework and Management Information Services, to the extent which these are defined.

NMP is implemented in layers 4 through 7 of the OSI model and is made independent of specific implementations of layers 1, 2, and 3 (see Figure 4). NMP is subject to some modification because it depends upon some OSI standards that are not yet finalized. The fol-

lowing section briefly describes the OSI layers and protocols in which NMP is implemented; the status of each OSI component is listed to provide a general picture of the degree to which NMP is subject to change.

OSI Layer 4—International Standard Status (finalized). The Transport Layer establishes (node to node) network connections, provides end-to-end data acknowledgment, terminates network connections, and performs related tasks. NMP's Transport Protocol and Transport Services for NMP are published in document TR54004.

OSI Layer 5—International Standard Status (finalized). The Session Layer establishes endpoint-to-endpoint sessions, specifies duplex or half-duplex service between session users, and coordinates session termination. Session Services for NMP and the Session Protocol for NMP are published in document TR54004.

OSI Layer 6—International Standard Status (finalized). The Presentation Layer ensures compatibility between incoming file, record, and data formats and the format requirements of the receiving systems. Presentation Services for NMP and the Presentation Protocol for NMP are published in document TR54004.

AT&T Unified Network Management Architecture (UNMA)

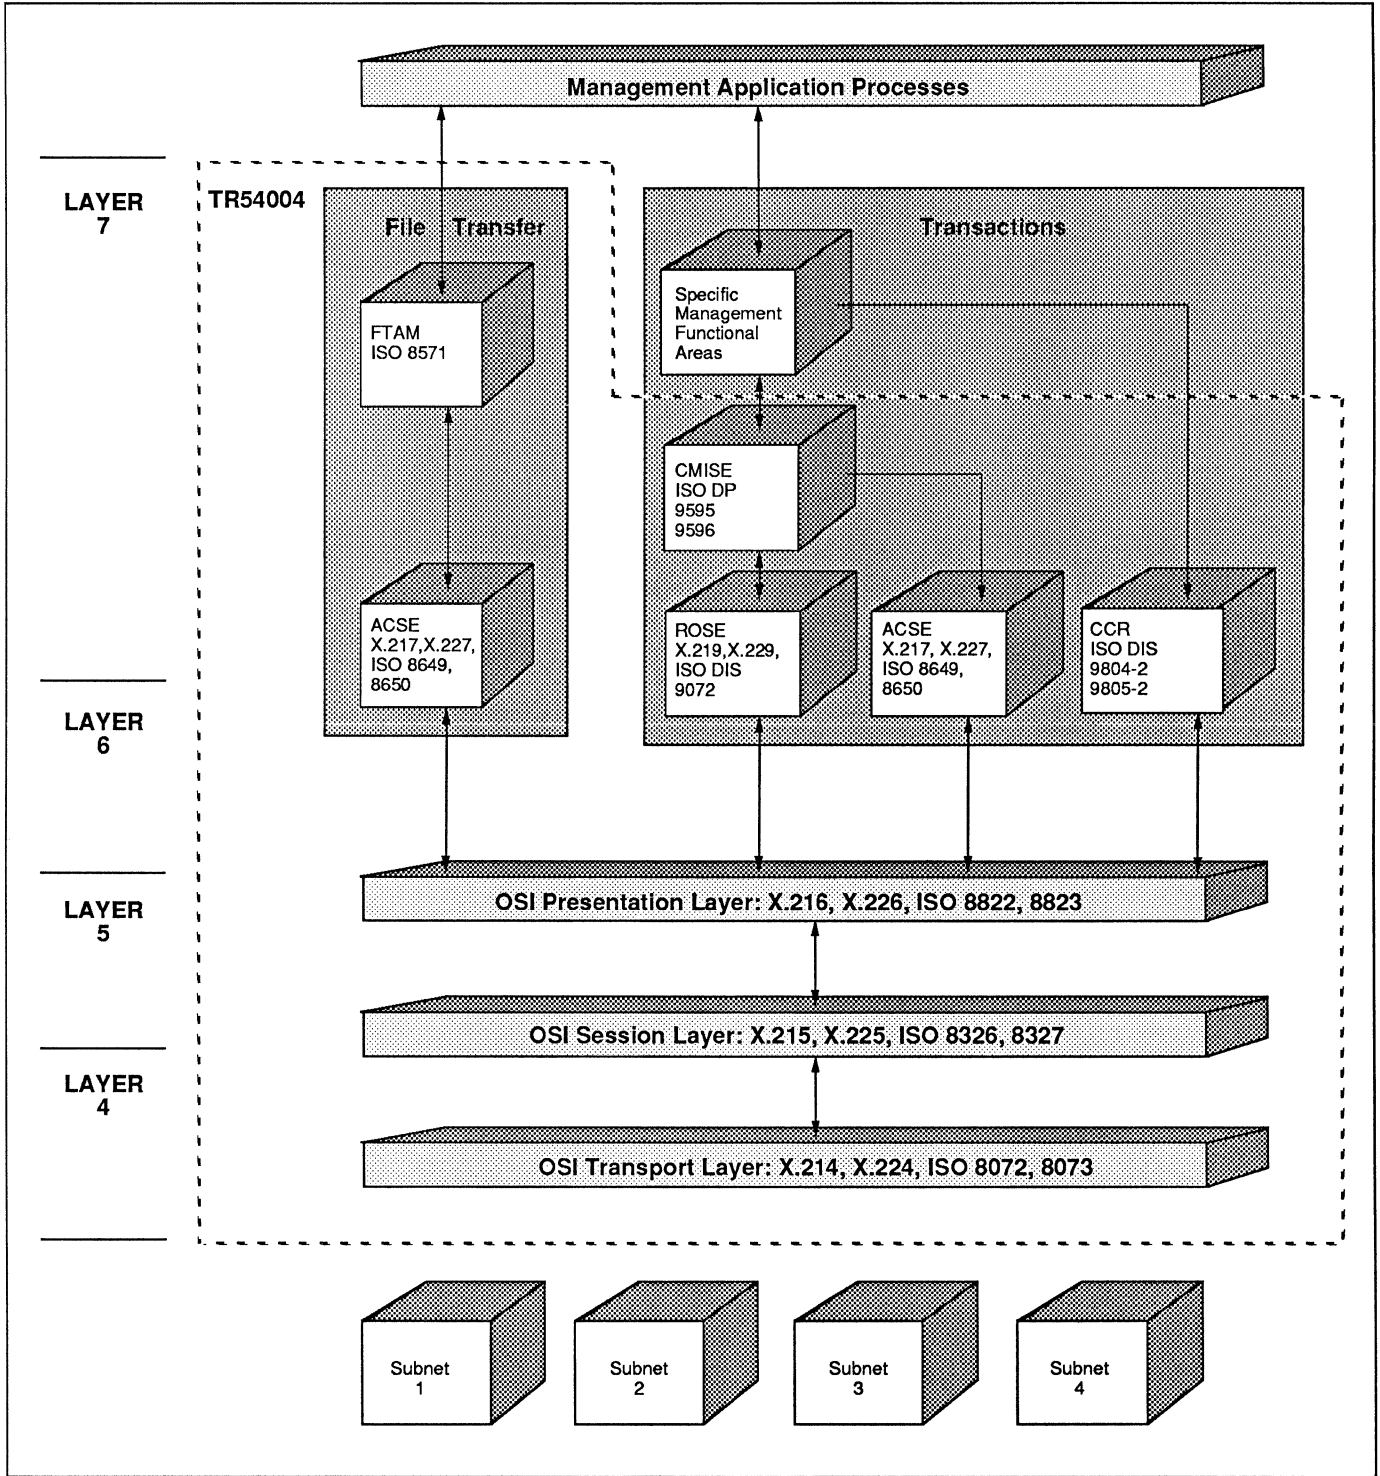


Figure 4. The subset of OSI standards and protocols contained in AT&T's NMP. AT&T document TR54004 (NMP Specification—Transport through Application Layers) describes the NMP implementation of OSI standards within the dotted lines on this illustration. AT&T published NMP specifications covering Specific Management Functional Areas (at layer 7), described in documents TR54005 and TR54006.

AT&T Unified Network Management Architecture (UNMA)

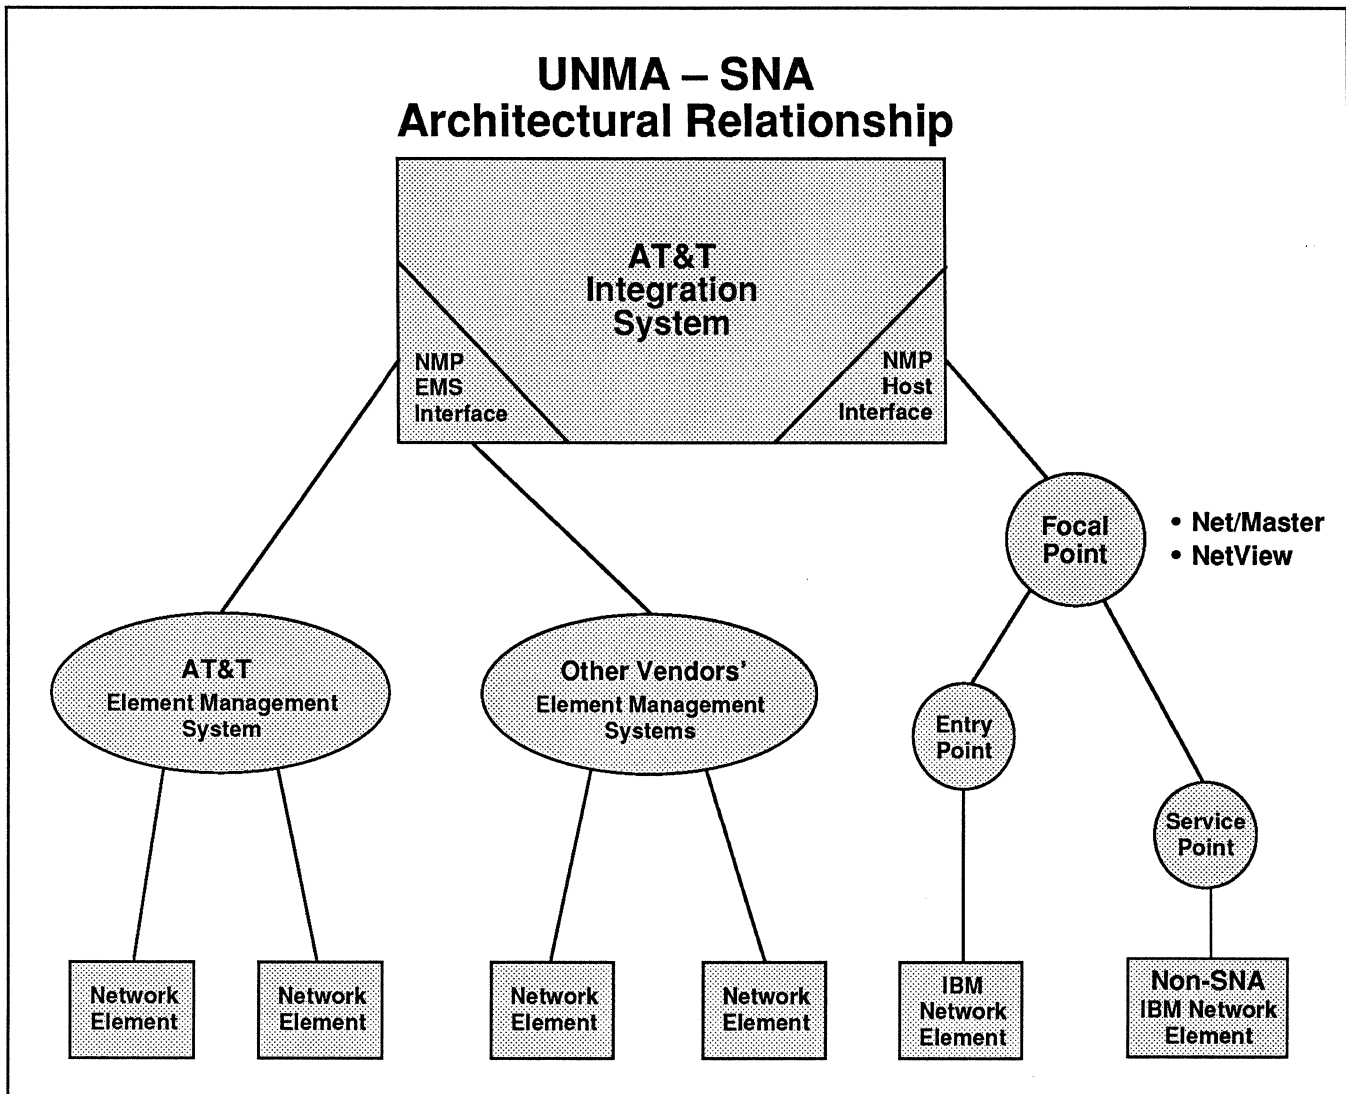


Figure 5. The relationship between the ACCUMASTER Integrator and the SNA environment. In the initial implementation, the ACCUMASTER Integrator will communicate to other vendors' network management systems (EMSs) via AT&T-supplied gateways. AT&T and Cincom Systems have jointly developed software (the AT&T-marketed SNA Management Application and the Cincom-marketed UNMA Application) that feeds the Integrator with logical data from the SNA focal point program (IBM's NetView or Cincom's Net/Master).

OSI Layer 7—Portions have attained International Standard Status (partially finalized). Within the Application Layer are several sublayers used for Network Management, including:

- **ASCE** (Association Control Service Elements)—International Standard (finalized)
- **FTAM** (File Transfer, Access, and Management)—International Standard (finalized)
- **ROSE** (Remote Operations Service Elements)—Draft International Standard (relatively stable)

- **CCR** (Commitment, Concurrency, and Recovery)—Draft International Standard (relatively stable)
- **CMIS** (Common Management Information Service)—Draft International Standard (relatively stable)
- **CMIP** (Common Management Information Protocol)—Draft International Standard (relatively stable)

AT&T divides its application layer services into two categories: transaction services and file transfer services. Transaction services are dependent upon CMIS,

AT&T Unified Network Management Architecture (UNMA)




Figure 6. AT&T's UNMA targets nine areas of functionality. AT&T has added the functions of integrated control, programmability, operations, and planning to the existing five-part OSI Functional Model.

ROSE, and ACSE. Enhanced transaction services, which provide for two-phase commitment and chaining and similar sophisticated facilities, depend upon CCR in addition to CMIS, ROSE, and ACSE. File transfer services require ACSE and FTAM. AT&T's implementation of these for NMP, as well as the CMIP protocol implementation, are outlined in document TR54004.

In addition to NMP's implementation across layers 4 through 7, AT&T has also published NMP application message sets for configuration management and fault management. (See the section entitled "UNMA Functions" for a description of configuration management, fault management, and the seven other UNMA functions.) As illustrated in Table 1, OSI specifications covering these two functions are closer to final approval than are specifications for performance, security, and accounting management. AT&T is currently working on message sets for the latter as well as for the other remaining UNMA network management functions, although progress to some extent depends on OSI committee progress. AT&T, however, sees no reason to wait for these functions to attain draft international standard (DIS) status before embarking on implementation. The company is now developing products that conform to OSI draft proposals (DPs), even though DPs are subject to modification. AT&T pledges to modify its products to comply with any OSI changes, however, and it expects those changes to be minor. Despite this aggressive approach to product development, it will take a number of years for AT&T to develop full functionality in all its targeted functional areas.

AT&T has pledged that it will modify NMP to conform to final OSI specifications. Notwithstanding, users must realize that the OSI guidelines leave room for

differences in vendor implementations—differences that can tarnish the lure of interoperability. Within each OSI layer are not only mandatory services but optional ones as well, which a vendor may choose to implement. In addition, vendors may choose to extend protocol definitions. For example, AT&T adds a parameter constraint in the ACSE Protocol for NMP that is not required in the ISO standard. Thus, users must be aware that while OSI-based implementations are *open*, this does not mean that they are *identical*.

Vendor Support for NMP

AT&T proposes that other vendors use the published NMP specification and message sets to develop UNMA interfaces to their products. Multivendor connectivity under UNMA depends upon vendor acceptance of NMP. Digital Equipment Corporation and a few other pioneers are also developing OSI implementations. Equipment vendors do not have unlimited resources, and many will be forced to choose between AT&T's NMP and other alternatives. Datapro believes that AT&T's network expertise and its tradition of superior research and development make NMP the best *currently available* technological alternative for vendors seeking an OSI-based approach. However, AT&T must win the support of the Bell Operating Companies and other vendors to make its private/public network integration plan viable. Furthermore, AT&T will need to deliver its UNMA solution with marketing savvy that it has heretofore lacked in the data world. This raises the question of whether choosing NMP is truly a safe bet.

Up until the second quarter of 1988, industry analysts viewed the lack of vendor support for NMP as a major hindrance to user acceptance of UNMA. AT&T has recently pledged to take a more proactive role in promoting NMP among equipment vendors for LECs and PTTs. AT&T currently lists 17 equipment vendors supporting NMP (see Table 2.) This list will undoubtedly grow between now and early 1990, when the OSI/Network Management (NM) Forum plans to stage an interoperability demonstration. (See "The OSI/NM Forum," below.)

The OSI/NM Forum

AT&T played an instrumental part in the July 1988 formation of the OSI/Network Management (NM) Forum. The other founding members were Hewlett-Packard, Unisys, Amdahl, British Telecom, Northern Telecom, Telecom Canada, and STC PLC (UK). Forum membership has increased to 59, including 6 new voting members and 45 associate members (see Table 3).

AT&T Unified Network Management Architecture (UNMA)

OSI MANAGEMENT STANDARDS			EXPECTED REGISTRATION DATES		
Title	Reference Document	Working Document	Draft Proposal	Draft Int'l. Standard	Int'l. Standard
OSI Management Framework	ISO 7498-4	complete	complete	complete	complete
OSI Management Information Service Overview	N 2683	complete	December 88	July 89	July 90
Structure of Management Information	N 2684	complete	complete	May 89	July 90
Common Management Information Service (CMIS)	ISO 9595	complete	complete	complete	September 89
Common Management Information Protocol (CMIP)	ISO 9596	complete	complete	complete	September 89
Configuration Management	N 2686	complete	complete	July 89	July 90
Fault Management	N 2687	complete	complete	July 89	July 90
Security Management	N 2688	complete	September 89	April 90	July
Accounting Management	N 2689	complete	April 90	September 90	April 91
Performance Management	N 2673	complete	September 89	April 90	April 91
Software Distribution	N 2671	December 88	September 89	April 90	April 91

Table 1. Expected registration dates for key OSI Management standards. Of the five network management functions, Configuration Management and Fault Management are closest to final text. More importantly, Common Management Information Service (CMIS) and Common Management Information Protocol (CMIP) are expected to attain International Standard status in September 1989.

The OSI/NM Forum's stated goal is to accelerate delivery of OSI Management-based products by forming a consensus on protocol options, message sets, and management object definitions. Forum members anticipate that such a consensus will not only promote product implementation development but also influence international bodies to adopt standards supporting (or, at least not conflicting with) those implementations. Forum members have pledged to promote Forum-adopted platforms within the international standards bodies.

The Forum's efforts will culminate in an interoperability demonstration, planned for mid-1990 (see Table 4). The Forum is currently in the process of defining the messages that will allow that demonstration to take place. Specifically, these messages address *event management*, which includes both fault management and configuration management. The Forum has already drafted an Application Services specification covering event management, which members are now reviewing. The closure date for a Forum consensus on messages is planned for mid- to late 1989. The Forum's first document, the OSI/NM Forum Protocol Specification, was approved by members in January 1989. The document is available for purchase and may be ordered by calling the Forum at (201) 766-1544. The protocol stack is fairly close to AT&T's published definition of NMP.

In parallel with finalizing its event management specifications, the Forum will be working toward a consensus on management architecture, including object definitions that name and define management data.

AT&T states that it will modify NMP to comply with whatever the OSI NM/Forum adopts and, ultimately, with whatever OSI committees adopt. Clearly, each vendor represented in the OSI/NM Forum has its own reasons for joining and will vote for an implementation that promotes those reasons. That is to say, there is no guarantee that OSI/NM Forum members will automatically adopt NMP message sets without modifications. On the other hand, regarding anticipated OSI/NM Forum decisions, AT&T has stated that "the technical work done on NMP will not change." AT&T is using a tremendous amount of resources for developing NMP into an appealing, viable protocol option. The company therefore anticipates that the OSI/NM Forum will adopt NMP message sets without too many changes.

The Forum has established the first quarter of 1990 as a target date for staging an OSI interoperability demonstration. Consequently, the next six months will clarify the force and focus of vendor acceptance of NMP, as exhibited in the OSI NM/Forum's actions and membership role.

In the interim, Digital Equipment Corporation is moving full speed ahead in garnering support for its own

AT&T Unified Network Management Architecture (UNMA)

Avante-Garde Computing, Inc.
Avanti Communications Corp.
Coastcom T1 Networking Products
Digital Communications Associates/Cohesive
Dynatech, Inc.
Emcom Corp.
General DataComm, Inc.
Hekimian Laboratories, Inc.
Infotron
Integrated Telecom Corp.
Kaptronix, Inc.
Newbridge Networks, Inc.
Paradyne Corp.
Racal-Milgo
Sync Research, Inc.
Telinq Systems
Tridom Corp.

Table 2. Vendors supporting AT&T's Network Management Protocol (NMP).

OSI-based approach, Enterprise Management Architecture (EMA). Announced in Cannes, France in September 1988, EMA employs a flexible director-entity approach that allows multiple *directors* (Digital network management systems) to manage *entities* (analogous to Network Elements) on multiple domains. EMA specifies a standard OSI-based format for director-to-director information exchange. At this time, however, Digital has not announced peer-to-peer exchange with its directors and NetView or with the ACCUMASTER Integrator. (For more information on EMA, see Report NM40-325-101, "Digital Equipment Corporation Enterprise Management Architecture.")

According to its press releases, Digital is committed to delivering OSI-based management products during 1990. By accomplishing this, Digital could damage AT&T's bid to wield primary influence over the direction of OSI-based implementations. If, however, AT&T's supporters continue to increase in number and if Digital announces no viable strategy for linking to NetView, that damage will probably be minimal, at least for the next few years.

CINCOM'S ROLE IN UNMA

In June 1988, AT&T and Cincom Systems, Inc. announced an agreement to jointly develop an application that will provide logical data to the UNMA integrating system. Cincom currently produces Net/Master, widely accredited as the only real competition to IBM's NetView. (For more information on Net/Master, see *Datapro Reports on Communications Software*, Report CMS60-153-101 "Cincom Systems, Inc. Net/Master.")

Cincom brings considerable SNA expertise to the UNMA development effort. The Cincom-developed software, called the UNMA Application, resides in the IBM host and allows the ACCUMASTER Integrator

VOTING MEMBERS:

Amdahl Corp.
AT&T
British Telecom
Digital Communications Associates, Inc.
GEC Plessey Telecommunications, Ltd.
Hewlett-Packard
MCI Telecommunications
Microtel, Ltd.
Nippon Telegraph and Telephone
Northern Telecom, Inc.
Societa Finanziaria Telefonica S.p.A.
STC plc
Telecom Canada
Unisys (Timeplex)

ASSOCIATE MEMBERS:

Alcatel, N.V.
Alantic Research Corp.
Avant-Garde Computing, Inc.
Bull S.A.
Cable and Wireless plc
Case Communications, Ltd.
CNCI Telecommunications
Computrol
Contel Technology Center
Data General Corp.
Dynatech Communications
Ericsson Business Communications AB
Fujitsu America, Inc.
Gartner Group
Hekimian Laboratories, Inc.
Hitachi Telecom (USA), Inc.
Infotron Systems Corp.
Interlan, Inc.
Kokusai Denshin Denwa Co., Ltd.
NCR Corp.
NEC America, Inc.
Network Equipment Technologies
Newbridge Networks Corp.
Nixdorf Computer Engineering Corp.
Novell, Inc.
OKI Electric Industry, Co., Ltd.
Paradyne Corp.
Philips Data Systems
Prime Computer, Inc.
Protocols Standards and Communications (PSC), Inc.
Racal-Milgo
Retix
Siemens AG
Sirti, S.P.A.
Tech-Nel Data Products, Ltd.
Telenet Communications Corp.
Telindus N.V.
Telwatch
Vance Systems, Inc.
Zellweger Telecommunications

Table 3. OSI Network Management/Forum members, as of May, 1989.

to interface with either IBM's NetView or Cincom's Net/Master. This application supports SNA customers by extracting logical information from IBM systems and feeding it to the SNA Management Application in the ACCUMASTER Integrator (see Figure 5). The ap-

AT&T Unified Network Management Architecture (UNMA)

UNMA TIMELINE							
UNMA Announced	1st NMP Spec Published	3 NMP Specs Published	AT&T Joins OSI NM/Forum	OSI NM/Forum Announces 26 New Members	Three Major Product Announcements: ACCUMASTER Integrator	OSI NM/Forum Approval of Protocol Spec	OSI NM/Forum Interoperability Demonstration Planned
	NetPartner Product Announced	SNA Link to UNMA Announced (Cincom Agreement)			ACCUMASTER Management Services Accumaster Consolidated Workstation (Release 2)		
4Q87	1Q88	2Q88	3Q88	4Q88	1Q89	2Q89	1Q90

Table 4. This timeline depicts milestones in the evolution of UNMA. In January 1989, AT&T announced its flagship product, the ACCUMASTER Integrator. This product is expected to be released in September 1989.

proach could feasibly provide peer-to-peer interfaces with other logical network management systems, such as Digital's EMA.

According to Cincom, the UNMA Application builds upon the functionality that Net/Master currently provides, but it will not preclude UNMA customers from using IBM's NetView. The Net/Master Advanced Network Management component, which retails for about \$15,000, is a prerequisite for using the UNMA Application, however. The UNMA Application, which costs approximately \$50,000, is itself a prerequisite for using the ACCUMASTER Integrator (see below).

According to Cincom, the UNMA Application does not depend on IBM's Network Performance Monitor (NPM). The current Net/Master product requires access to the NPM, which IBM could conceivably bundle in with NetView at a future date. (There have been no indications so far that IBM intends to do that, however, and Datapro believes that it is unlikely.)

UNMA PRODUCTS AND SERVICES

The ACCUMASTER Integrator

The ACCUMASTER Integrator is a customer premise product that was announced on January 31, 1989. It is now in beta test and will be available to customers in September 1989. The Integrator is a UNIX system that runs on an AT&T 3B2/600 in its initial implementation. AT&T plans to offer additional hardware platforms in the future—such as a lower cost implementation for smaller customers, and for larger

customers, another platform offering more power to accommodate added functionality. The Integrator includes a Sun Workstation that displays color-graphic network views. The operators can use a mouse to manipulate icons that depict network components.

Initially, the Integrator will provide control in two areas—alarm integration and configuration management. AT&T plans to expand this functionality in the future. The Integrator also features *alarm correlation*—a capability not yet present in NetView. Alarm correlation compares alarms received by the Integrator, determines the most likely cause of the alarm, and suppresses secondary alarms linked to the original problem. The Integrator uses network device profiles (stored in its internal Informix database) in combination with its own internal logic to perform alarm correlation. This feature may help operators determine what to do first and whether the alarms are related by isolating the most likely source of the alarm. However, the process is not infallible, since it can determine only the *most likely* source, not the *actual proven source*.

The Integrator product solicits *logical* network management information from the UNMA Application and combines it with *physical* network data transmitted (via an NMP interface) from various Element Management Systems. The first release of the ACCUMASTER Integrator features a short-stack implementation capable of automatically uploading changes from EMSs. For those equipment vendors that have not yet implemented the NMP interface on their products, AT&T will supply gateways to make it easier to interface with the Integrator.

AT&T Unified Network Management Architecture (UNMA)

ELEMENT MANAGEMENT SYSTEM	NETWORK ELEMENT
Data Systems	
DATAPHONE II System Controller DATAPHONE II ACCULINK Network Manager STARKEEPER Network Management System StarLAN Software UNIX Software 6544 Software Customer Test Service	Modems, multiplexers Modems, multiplexers DATAKIT VCS, ISN StarLAN 3B2 Computers 6500 MCS Dataphone Digital Service
Voice Systems	
Trouble Tracker Centralized System Management Centralized System Management, VMAAP Multi-Function Operations (MFO) System Customer Network Control Center Service Management System/MISR Routing Control Service	System 85, 75, Dimension PBXs System 75 System 85, Dimension 5ESS Switches, PBXs EPSCS SDN Advanced 800 Service
Voice and Data Systems	
Customer-Controlled Reconfiguration Order Management Service	ACCUNET T1.5 Provisioning

Table 5. AT&T Element Management System (EMS) offerings.

In addition to processing management data through the NMP interface, the ACCUMASTER Integrator includes terminal emulation capabilities that enable the user to cut through to the individual Element Management Systems. This important feature gives the user the full effect of operating at the Element Management System console—such as a STARKEEPER console. This provides an advantage over the NetView/PC implementation, which presents generic alerts or character strings (extracted from the other physical network devices using code points) rather than the actual EMS console display.

To provide the cut-through capability, AT&T includes a 3279 emulation package as part of the Integrator. The Integrator console (Sun Workstation) has the capability of displaying information just as it appears on the NetView console. The Integrator console becomes a terminal defined in SNA and, as such, it can access SNA applications such as Cincom's Net/Master or IBM's NetView.

The ACCUMASTER Integrator facilitates an end-to-end view by creating a unified, "virtual network" composed of both private and public network elements. The integrating system accomplishes this by combining information from customer premise Element Management Systems (EMSs) with information from EMSs controlling those portions of the public network partially allocated to the customer. Under UNMA, this integrating system may reside either on the customer premises or in the public network, or in both simultaneously.

A typical Integrator configuration supporting several Element Management Systems costs approximately \$300,000. This cost does not include Cincom's UNMA Application (\$50,000) and the required Net/Master Advanced Network Management component (about \$15,000).

NetPartner Network Management System

AT&T's integrating system for LECs is the NetPartner Network Management System (NMS). NetPartner NMS is available now from AT&T. This product is a hardware/software combination that gives Centrex customers restricted access to operations systems embedded within the phone company's network. Specifically, NetPartner translates information that is currently within the phone company's operations systems into a form that is easier for customers to understand, manipulate, track, and alter. The phone company maintains control over what the customer can access.

NetPartner NMS features a three-part architecture, composed of host equipment at the phone company, customer premise equipment, and operations systems at various phone company sites. The equipment requirements are as follows:

- **Customer Premise Equipment**—One SUN-3 Series Workstation or an AT&T MS-DOS-based personal computer. (The SUN Workstation is required if windowing and graphics capabilities are desired.) Re-

AT&T Unified Network Management Architecture (UNMA)

mote operations systems must be connected via private lines or must be accessible to a DATAKIT VCS.

- **Phone Company Equipment**—Three AT&T 3B2/600 host computers, a communications processor, and a workstation for administration, as well as a DATAKIT Virtual Circuit Switch are required. A printer and voice synthesizer for audible alarms are optional.
- **Phone Company Operations Systems**—NetPartner provides access to the following phone company operations:

Engineering and Administrative Data Acquisition System (EADAS)—collects and analyzes traffic over the telco switch, assisting customers to determine their capacity requirements.

Switching Control Center System (SCCS)—centralizes switch and network terminal maintenance; helps the customer determine responsibility for telecommunications failures.

Loop Maintenance Operations System (LMOS) and Mechanized Loop Testing (MLT)—will allow customers to see the status of trouble tickets associated with their lines.

Switch Applications Processor (1A ESS or 5 ESS)—allows customers to request Station Message Detail Recording reports to monitor outgoing call traffic.

MacStar End Customer Management System—allows customers to manipulate Centrex and ISDN lines and route selection, among other services.

In addition to these five phone company operations, NetPartner provides a 3270 terminal emulation capability that allows access to IBM's NetView or NCCF/NPDA or Cincom's Net/Master. In this configuration, NetPartner appears like a cluster controller to the IBM host computer.

The NetPartner system provides a menu-driven, graphics-based user interface. The phone company customizes each menu to show only those functions that the customer's company has purchased. Currently, one NetPartner Network Management System can support 5 to 10 customers. The phone company defines the maximum number of users per customer; however, no more than 20 users can access NetPartner at one time. Future NetPartner enhancements are planned, including an ISO-based interface to specific customer premises-based and interexchange carrier products. Some artificial intelligence-based capability is also planned.

AT&T is directing its NetPartner marketing efforts at phone companies and end users. AT&T is seeking to convince LECs that end users are willing to pay extra for the control facilities that a NetPartner-equipped LEC can offer. AT&T is trying to persuade end users to demand NetPartner capability from their local telcos. AT&T is particularly pursuing those users interested in ISDN capabilities. Pricing information on NetPartner is not yet available.

Network Operations Center

In addition to the ACCUMASTER Integrator and NetPartner, AT&T has announced an **integrating service** called ACCUMASTER Management Services. This service is administered from various Network Operations Centers (NOCs) that physically reside either in AT&T's network or on the customer's premises. NOCs will be staffed by AT&T personnel teams solely dedicated to a specific customer. When under contract for NOC services, the customer will receive end-to-end provisioning and maintenance covering customer premise equipment and services supplied by other vendors, as well as AT&T. Provisioning will include station moves, changes, and relocations; individual PBX, DCE, and CKT orders; and facility and premises provisioning. In addition, call receipt, trouble tracking, repair verification, restoration planning, and implementation services will be provided. In the future, NOCs will offer additional services beyond these to ACCUMASTER Management Services customers.

There are roughly 10 NOCs already in service. About half of these are on the customer's premises. AT&T has not released pricing information on NOC/ACCUMASTER Management Services.

ACCUMASTER Consolidated Workstation

The AT&T ACCUMASTER Consolidated Workstation (ACW) was announced in September 1987 and became available on mid-1988. Release 2 of the ACW, announced in January 31, 1989, runs on an AT&T 6286 or 6312 Work Group System (WGS) and uses multitasking applications software to monitor multiple systems simultaneously. The ACW requires MS-DOS 3.2 or 3.3 and Microsoft Windows Version 2.1.

The ACW displays a consolidated (not integrated) network view by providing separate windowing sessions to various AT&T Element Management Systems. The ACCUMASTER Consolidated Workstation currently supports the following AT&T element management systems:

AT&T Unified Network Management Architecture (UNMA)

- ACCUMASTER Trouble Tracker, VMAAP
- ACCUNET T1.5 Service with Customer Controlled Reconfiguration (CCR)
- DATAPHONE II System Controller
- DATAPHONE II ACCULINK Network Manager
- DATAPHONE II 839A Dial Backup System
- STARKEEPER Network Management System

ACW users can open window sessions to any four of these systems simultaneously.

In addition, the ACW's 3270 terminal emulation facility provides IBM host access from NetView or Cincom's Net/Master. Users can communicate with host applications while maintaining active sessions with other Element Management Systems. ACW Release 2 also includes both VT100 terminal emulation and 513 emulation. The 513 emulation feature gives users the capability to access the AT&T System 75/85's (PBX) Centralized System Management and other vendors' Element Management Systems. The VT100 capability enables users to access the AT&T network-based information, such as Management Information Systems Report (MISR). MISR is a network administration report system designed to support Software Defined Networks (SDNs).

ACW operators can observe network changes in real-time and interact with the appropriate system to perform testing, problem diagnosis, troubleshooting, and reconfiguration.

CURRENTLY AVAILABLE AT&T ELEMENT MANAGEMENT SYSTEMS

There are about 20 customer-accessible AT&T Element Management Systems. An additional 100 network-based EMSs deployed by AT&T are not currently customer accessible. Over the next decade, AT&T plans to make more of these systems accessible, providing customers with additional trouble ticketing information, as well as alarm and performance information on AT&T facilities.

All currently customer-accessible EMSs are listed in Table 5.

UNMA FUNCTIONS

Network management systems are necessary in order to obtain realtime information on network perfor-

mance and traffic characteristics, diagnose problems, and reconfigure to meet changing needs. In the past, network management was characterized by separate management systems devoted to providing these services for a particular vendor's product or group of products. The AT&T Element Management Systems described in the previous section are typical examples.

With UNMA, AT&T proposes to integrate data from disparate management systems and collectively provide critical network management functions. Specifically, UNMA defines nine functions to support users in managing their networks (see Figure 6). The first five of these functions are included in the OSI Functional Model of management. The generic definitions of these functions are followed by specific UNMA facilities described in italicized print:

- **Fault Management**—The goal of fault management is to maintain network availability at an acceptable level. On a day-to-day basis, this means quick and accurate problem detection and problem determination.

The ACCUMASTER Integrator correlates alarm information from various systems to pinpoint the event or fault that may have caused multiple alarms. UNMA calls for complete audit trail of the fault management process, supported by trouble tracking system.

In April 1988, AT&T published the Network Management Protocol Fault Management Message Set Specification. This document outlines how NMP enables users to control alarm reporting, manipulate alarm information, and set alarm reporting parameters.

- **Configuration Management**—The goal of configuration management is to manipulate network configurations to adapt to changing needs and traffic patterns or to isolate problems. To support this, systems must collect information on the current state of the network, noting changes; modify network attributes; and change configuration.

The UNMA definition of this function has four basic aspects: network inventory management, change management and provisioning, name management, and actual connections (relating inventory items to their physical layout.) Network inventory management involves tracking all devices, services, and systems on the network. Change management and provisioning supports both scheduled and unscheduled movement of telephones, modems, terminals, circuits, and other network components. Name management governs the network directory. In April 1988, AT&T published the NMP Configuration Management Message Set Specification. This document

AT&T Unified Network Management Architecture (UNMA)

outlines how NMP enables users to control configuration management services and access/change resource configuration information.

- **Performance Management**—The goal of performance management is to identify and correct potential problems before they cause a fault. To accomplish this, the management system must collect data on current network and resource performance levels and maintain performance logs.

UNMA provides the capability to correlate information from multiple systems to help the users identify network performance trends. UNMA users will have the capability to monitor selected network components via user-definable measures and thresholds. UNMA will provide both recent history and current performance data, which the user can analyze in order to identify network performance trends.

- **Accounting Management**—This function informs users of costs incurred and enables users to set accounting limits.

Under UNMA, users can compare usage and billing information across related systems and obtain more complete billing, verification, and chargeback information. Chargebacks may include charges for fixed-cost items such as telephones or terminals as well as usage; UNMA supports comparison of vendor bill verification and comparison of vendor bills with inventory and internal measures.

- **Security Management**—Security management encompasses access control, authorization facilities, and partitioning the network. The OSI definition of security management also includes support for encryption and key management and the maintenance and manipulation of security logs.

UNMA supports tracking logon attempts and violations to prevent unauthorized network access. UNMA also supports multiple network management permission levels. Under UNMA, network administrators can manage the network from either one or several different network management operation centers by partitioning the network.

- **Planning**—While not defined as an OSI Management function, planning is widely accepted as a major network management subsystem. The goal of planning is to design and optimize models that describe potential changes to the network. Users must consolidate usage trends and performance data to support this. Planning typically involves collecting

performance data, consolidating usage trends, and analyzing future requirements.

UNMA outlines three common types of planning: capacity planning (day-to-day fine tuning, such as adding or rearranging trunks); contingency planning (backup and disaster recovery, including estimated costs); and strategic planning (new applications, growth plans, acquisitions, reorganizations).

- **Operations Support**—This encompasses managing the staffing and operation of a network management center.

UNMA defines four aspects of operations support: creating network management center procedures (for trouble logs, maintenance fixes, shift changes, etc.); analyzing work and information flow at the center; analyzing network management center staff requirements; and preparing user training and development plans. AT&T's Management Services provides operations support.

- **Programmability**—The goal of programmability is to customize the network management system to meet corporate needs. There are no off-the-shelf solutions in network management. Programmability is critical because each network is different and networks are characteristically in a state of flux.

UNMA provides parameterization of key system characteristics, flexible report capabilities, customizable scripts, and custom programming options. UNMA supports programmability in C language.

- **Integrated Control**—The goal of integrated control is to create the image of a single virtual network, even though it may actually comprise diverse, separate management systems.

AT&T currently provides integrated control with its ACCUMASTER Integrator.

These nine generic function descriptions provide a useful framework for evaluating integrated network management systems. Currently there is no single OSI-based product or service that provides comprehensive, integrated support in all of these functional areas. Over the next decade, AT&T plans to evolve UNMA and its ACCUMASTER product line to fill this gap. If AT&T successfully turns these plans into products in a timely fashion, UNMA could turn out to be "the right choice." □

Digital Equipment Corporation Enterprise Management Architecture (EMA)

This report will help you to:

- Compare EMA's distributed approach to NetView's strictly hierarchical framework.
 - Assess the viability of Digital's approach to network management.
 - Decide whether your organization's network could benefit by using Digital's flexible management architecture.
-

Enterprise Management Architecture (EMA) is Digital Equipment Corporation's strategic plan for providing integrated network, system, and application management. EMA provides a highly flexible means for managing customer networks—whether those networks exhibit a centralized hierarchical structure or a flat, distributed structure—and most anything in between.

EMA will incorporate an implementation of the Common Management Information Protocol (CMIP) to facilitate communications between EMA software and third-party systems. CMIP is OSI's protocol for exchanging management information among systems within a multivendor environment.

In comparison, IBM's NetView lacks meaningful OSI support and cannot yet effectively manage distributed networks. While AT&T's UNMA provides an OSI-like structure, it is not quite as flexible as EMA. AT&T is much closer than Digital to delivering a usable product, however.

This report describes EMA and evaluates some of its strengths and weaknesses with respect to its two main competitors—IBM's NetView and AT&T's UNMA/ACCUMASTER Integrator.

In September 1988, Digital Equipment Corporation announced its platform for integrated network management—Enterprise Management Architecture (EMA). For several years, Digital has provided sepa-

rate network monitoring and control products for DECnet systems and Ethernet local area networks (LANs). EMA is significant in that it is Digital's first major attempt at integrating existing product capabilities and setting a strategy for long-term network management product evolution.

Analysts and users anticipated Digital's network management announcement long before EMA was officially unveiled. Digital's major rivals, notably IBM and AT&T, introduced integrated network management strategies (and, in IBM's case, actual products) one to two years earlier. In the spring of 1986, IBM became the first major vendor to provide an integrated network management package (NetView). AT&T announced its Unified Network Management (UNMA) architecture in September 1987 and its participation as a founding member of the OSI Network Management/Forum in July 1988. These events set

REPORT HIGHLIGHTS:	PAGE
HOW EMA ADDRESSES THE MARKET	102
EMA'S STRUCTURE	103
MANAGEMENT DOMAINS	105
FROM DECnet TO EMA	107

Digital Equipment Corporation Enterprise Management Architecture (EMA)

3Q88	2Q89	1Q90	3Q90
EMA Announced	Expected Publication Date of EMA Interface Specs	Expected Release of EMA Developer's Kit	Expected Release Date of CMIP-based Network Management for Digital's Products
Vendor Support for EMA Announced		Expected Release	Expected Release

Table 1. Digital's enterprise management architecture (EMA) timeline.

the stage for an expected announcement from Digital, which materialized two months later in September 1988.

Between the spring of 1986 and the fall of 1988, however, Digital was not idle. The company actually began developing its OSI-based architecture over two and one-half years ago. The task was enormous—develop a plan for migrating the world's largest installed networked base of 210,000 CPUs to OSI. To do this, Digital interviewed hundreds of its customers, asking them what they would like to see in an integrated network management system.

The result of this massive undertaking: a flexible, standards-based architecture that is highly adaptable to distributed computing environments. Digital has delivered the blueprint—but the actual products will probably not appear until 1990. (See Table 1.)

HOW EMA ADDRESSES THE MARKET

EMA incorporates a number of design goals and principles (which Digital calls "metrics"). As with IBM's NetView and AT&T's ACCUMASTER Integrator, EMA's primary function is to provide integrated network management—a consistent user interface, a common information repository, and integrated access to management functions and managed devices. However, four distinguishing EMA traits in particular directly address current market needs: EMA's applicability for distributed environments; EMA's open interfaces; third-party vendor support for EMA; and eventual OSI compliance (and a migration path to accomplish that end).

Distributed Processing: EMA's structure is exceptionally flexible and thus adaptable to distributed processing environments. (See the section entitled "Structure of EMA" for an explanation of why this is true.) Distributed processing and embedded management capabilities have been a part of Digital's philosophy for some time. EMA enhances Digital's embedded

management approach by providing a means of centrally coordinating the management of distributed systems.

Open Interfaces: The second strength, open interfaces, is really a given; ever since IBM published its Network Management Vector Transport (NMVT) interfaces in the spring of 1986, proprietary interfaces are not a viable option in network management architectures. Digital expects to publish its EMA interface specifications in the "EMA System Reference Manual" in the summer of 1989.

Third-Party Vendor Support: Digital seeks to offer EMA as a platform on which other vendors can build—a network management operating system that it will OEM to third parties. Digital had this approach in mind when, in developing EMA, it actively sought input from third parties as to what *they* would like to see Digital provide. This approach regarding third-party vendor relationships gives Digital several advantages over IBM. There are seven vendors currently working with Digital to develop the EMA interface specification. (See Table 2.) These vendors represent a balance of both voice and data products—thus, it would appear that the EMA interface specs should adequately accommodate requirements for both voice and data requirements.

<p>T1: Digital Communications Associates/Cohesive Stratacom, Inc. Timeplex, Inc.</p> <p>Modems: Codex Corporation</p> <p>PBX: Siemens AG Telecom Services Bureau International</p> <p>Bridges: Vitalink Communications Corporation</p>
--

Table 2. Vendors supporting digital's enterprise management architecture (EMA).

Digital Equipment Corporation Enterprise Management Architecture (EMA)

The seven vendors, which have pledged to develop EMA interfaces to their products, put a measure of clout behind Digital's platform and deliver yet another blow to IBM's weakening platform of vendor support for NetView/PC.

The EMA interfaces will support two-way monitoring and control between the EMA modules and the other vendors' products. Thus, on a technical level as well, Digital is courting other vendors in a style very different from IBM's. Digital's EMA will support peer-to-peer connections with other vendors' device management systems. IBM's NetView treats third-party systems as dumb terminals and severely restricts what other vendors can do.

AT&T is also actively pursuing vendor alliances for its UNMA/ACCUMASTER product line; most notably, it jointly developed an (strategic) ACCUMASTER component with Cincom Systems, Inc. This component, called the SNA Management Application, provides AT&T's entry product into managing SNA networks. AT&T also claims that about 20 vendors are committed to developing interfaces to its ACCUMASTER Integrator. AT&T has gone one step beyond Digital, however, in helping to form the OSI Network Management (NM)/Forum. The Forum, created to accelerate development and acceptance of OSI management products, will undoubtedly serve AT&T's interest by forming a consensus that closely resembles AT&T's CMIP implementation (called Network Management Protocol, or NMP). For more information about AT&T and the OSI NM/Forum, see Report NM40-313-101, "AT&T's Unified Network Management Architecture."

Eventual OSI Compliance: EMA's eventual OSI compliance and migration path contrast with IBM's approach to network management. When IBM developed NetView, it intended to *create* a de facto standard rather than *conform* to one. Moreover, Digital is offering a migration path to OSI. In September 1988, IBM announced that NetView will support OSI links. IBM's mechanism of support—using the CMIP protocol to communicate to the focal point—appears difficult to implement and reveals IBM's schizophrenic commitment to the goals of OSI. OSI is fundamentally a peer-to-peer model; NetView imposes a hierarchical structure on OSI management communications. NetView's planned support of CMIP does not assist in evolving user's networks to the distributed, heterogeneous model of connectivity embodied in OSI.

EMA's STRUCTURE

The structure of EMA is simple, yet flexible enough to be implemented in numerous ways. EMA is composed of several basic pieces; yet these pieces do not have to reside in one location. Rather, they can be broken up and distributed in any number of ways to reside on various systems or processors located throughout the network. This flexibility is a major advantage, since it is adaptable to the distributed computing environment which, due to the influx of PCs, LANs, and server technology, is on the rise in more and more organizations.

EMA employs a Director-Entity model to describe the relationship between the network components being managed (Entities) and the systems managing them (Directors). For those familiar with SNA, a Director fulfills the role of NetView, although as we shall see, a Director's structure and (implementation) is quite unlike NetView's.

ENTITIES AND DIRECTORS

An **Entity** is actually composed of two parts—the *managed object* (modem, communications line, etc.) and its *agent*—management software. This software acts as a conduit for management operations (events and directives) and may also provide a degree of management capability for the entity.

Under EMA guidelines, an agent and the managed object need not reside on the same system. For example, an agent may use a selected protocol over RS-232 lines to communicate with the actual managed objects—as long as the EMA Director has access to the exchanged information.

The **Director** is a software system that acts as an interface between the user and the managed network devices and systems (entities). It is composed of five parts: the Executive; the Management Information Repository (MIR); and three types of management modules—Presentation Modules, Function Modules, and Access Modules. (See Figure 1.)

While an Entity is typically embodied in a discrete product such as a modem, computer (DECnet node), LAN bridge, etc., a Director is more often a logical concept that is composed of multiple software components which may or may not reside in one physical location.

Digital Equipment Corporation Enterprise Management Architecture (EMA)

THE EXECUTIVE

The **Executive**, the core of EMA, maintains management information in a Management Information Repository. The Executive is a master control program written by Digital that provides standard interfaces and a common set of module support routines for coordinating the activities of the Presentation Modules, Function Modules, and Access Modules.

Chief among the Executive's module support routines are those supporting intermodule communications, implemented via a type of remote procedure call. This facility is critical, since EMA allows for multiple, distributed Directors communicating in a peer-to-peer fashion. A Director can request information from another Director's access module, or it may request services provided by another Director's Function Modules. Thus, a Director can use the services of modules that are distributed in various physical locations. Thus, Executive support for intermodule communications is critical to preserving the flexibility of EMA's design.

The Executive uses a dispatch mechanism to support intermodule calls. The dispatch mechanism determines the destination procedure, passes control to it, and returns response information back to the requesting procedure.

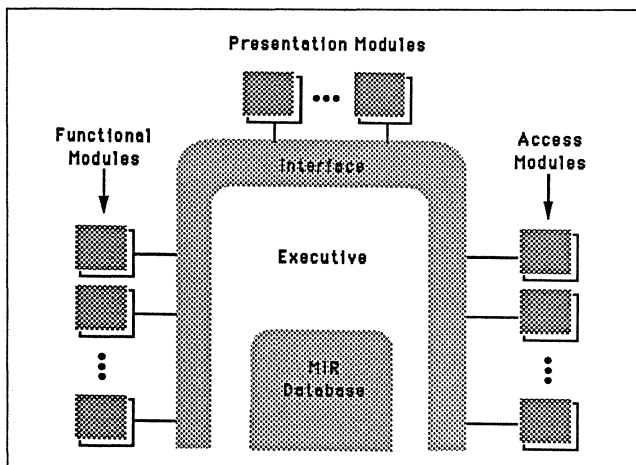


Figure 1. The structure of Digital's Enterprise Management Architecture (EMA) supports plug-in modules (Access Modules, Presentation Modules, and Functional Modules) that may be developed by Digital, third-party vendors, or by users themselves. At the heart of EMA's Executive control program is the Management Information Repository (MIR), an object-oriented database.

MANAGEMENT MODULES

Management Modules are the "plug in" components of a Director. These modules can interoperate without requiring specific predefined knowledge of each other; the Executive controls the intermodule communications process (described above) that is essential for the modules to interact. There are three types of modules: Presentation, Functional, and Access.

Presentation Modules create the user interface for EMA Directors. A Presentation Module can be written to support a specific console or presentation format or to interface with non-EMA applications. Digital designed EMA to allow for a number of Presentation Modules per Director, if the user so desires. For example, one Presentation Module might support a graphics workstation, and another might compile usage reports based on 4GL input. One advantage to this approach is the ability to support multiple different interfaces, independent of the network devices and subsystems (entities) which the Director manages.

It is expected that Digital will offer a standard set of Presentation Modules; however, the company will also encourage third parties to develop Presentation Modules for specific applications.

Functional Modules provide Configuration, Fault, Performance, Accounting, and Security Management—the five functional areas defined by the OSI Management standards. (For more information on these standards, see Report NM40-200-101, "OSI-Based Network Management.") EMA allows a Functional Module to encompass more than one function, however, and thus does not strictly define modules according to the five OSI categories. This is quite understandable, since the OSI committees themselves have had quite a difficult time drawing the line between such functions as Fault Management and Configuration Management, for example. An EMA Functional Module may, therefore, provide network design services using fault, performance, and configuration information. This measure of flexibility makes EMA highly adaptable not only for a user's current needs, but for future requirements as well—for example, when advances in expert system technology make it possible to design Functional Modules that perform high-level management tasks which correlate data from several OSI functional areas.

Functional Modules can be developed by Digital or by third parties wishing to build upon the EMA platform. It is possible that users may also want to develop their own customized Functional Modules in-house.

Digital Equipment Corporation Enterprise Management Architecture (EMA)

Access Modules communicate directly with typical network components such as PBXs, multiplexers, modems, and CSUs/DSUs and nonnetwork components such as systems and applications. To accomplish this, Access Modules implement protocols—thereby allowing the Directors which manage network components to remain protocol independent.

Access modules are functionally equivalent to programs which third-party vendors must write to interface with IBM's NetView/PC. EMA Access Modules provide some fundamental, positive differences, however. Access Modules translate specific device information into a data format understood by both Directors and other Management Modules (Functional, Presentation, and other Access Modules). Access Modules can, therefore, support two-way monitoring and control between Directors and managed devices—a feature that NetView/PC does not currently have. (NetView/PC Release 2 promises it, however.)

When an Access Module is attached to an EMA Director, the Access Module must be “enrolled.” During enrollment, the Access Module must inform its Director what it intends to manage and how. Specifically, it tells the Director, “These are the classes of entities (devices) I can access, and these are the management operations which those devices can support.” The Access Modules must then supply the Management Information Repository with certain information to be used by other Management Modules.

Access Modules are developed by third-party vendors that wish to offer EMA-compatible products. Digital recommends, but does not require, that these modules implement a CMIP-compliant protocol within the Access Module to enable the vendor's product to communicate with the EMA Director. Access Modules are the means by which EMA will support multivendor networks and, using CMIP-like protocols, will serve to achieve that goal. It is presumed that, upon the initial release of an EMA product, Digital will offer a set of Access Modules for managing Digital nodes.

Management Information Repository (MIR) is an object-oriented configuration database for information about network devices and management activities. According to Digital, the Repository is “independent of any particular implementation and is not limited to any specific database management system.” This provides a much more flexible approach than either IBM's NetView (which requires IBM's Virtual Storage Access Method [VSAM] and, as yet, has no comprehensive Repository) or AT&T's ACCUMASTER Integrator, which uses an Informix database.

The Management Information Repository acquires its stored information during the “enrollment” process—when an Access, Presentation, or Functional Module is attached to a Director. When a Module supplies the MIR with management information, other modules can then access that information through intermodule calls supported by the Executive's dispatch mechanism.

Four types of data are stored in the Management Information Repository: Class, Instance, Attribute, and Private. The first three types apply to managed devices and closely resemble OSI data types (and, consequently, the data types implemented in AT&T's UNMA.) The fourth type, Private Data, is typically associated with and is used by a single management module.

Class Data groups network devices (managed entities) into related classes. Bridges, terminal services, T1 processors, and DBMSs are examples of general entity classes. Entity Class Data's function is similar to a database data dictionary—management modules can query to find out a class entity's structure and the types of operations that may be performed on that class, for example.

Instance Data consists of configuration data, such as network addresses, node names, disk controllers, etc., that lie within a Director's sphere of control. (Actual managed network devices are referred to as Entity Instances.)

Attribute Data is specific management information, recorded and stored historically over time, that relates to managed network devices (entity instances). Attribute data is typically provided by an Access Module which collects the information either by polling or by receiving event messages (faults).

Private Data is a type of catchall category for data that relates to one particular Management Module (such as binary files of parse tables used by a Presentation Module). EMA imposes no uniform structure on Private Data, so it may be tailored to the specific requirements of the particular Management Module.

MANAGEMENT DOMAINS

Digital's EMA employs the concept of *management domains* to add yet another dimension of architecture flexibility. Digital defines a management domain as “a user-defined sphere of management interest and control.” This sphere is composed of at least one Director, along with network components and de-

Digital Equipment Corporation Enterprise Management Architecture (EMA)




Figure 2. A hierarchical, overlapping multidomain configuration. A corporate data center is linked via a packet switched network to three regional data centers. The regional centers, in turn, are linked via leased lines to branch office minicomputers. Branch office minis support PCs in various LAN configurations.

vices (entities) that share a common purpose—this purpose, or category, is purely arbitrary and determined solely by the user.

This contrasts sharply with the *domains* found in IBM's SNA. SNA domains are defined according to strict IBM specifications. (An SNA domain is a collection of Physical Units [PUs] and Logical Units [LUs] managed by a Systems Services Control Point [SSCP], a part of ACF/VTAM which *must* reside in an IBM mainframe.)

EMA domains are defined on the basis of *global entities*. A global entity is any network component or device that possesses a unique name or address. PCs, bridges, terminal servers, and nodes are examples of typical global entities. (Global entities do not belong

to any other parent class, but, rather, are the highest level of addressable entities within a network. Again, a user can choose to define a particular entity as global or as part of a subclass, keeping even the definition of a domain flexible.)

Domains may be defined according to:

- **Function**—grouping network components according to the function they perform, such as performance, fault, configuration, security, or accounting management.
- **Organization**—grouping network components according to corporate departments or subsidiaries that use them.

Digital Equipment Corporation Enterprise Management Architecture (EMA)

- **Technology**—grouping network components and devices according to their technological type, such as bridges, modem controllers, nodes, etc.
- **Geography**—bounded by DECnet area, Ethernet segment, building, time zone, etc.
- **Other**—any arbitrary grouping of the user's choosing.

Moreover, a domain may:

- Contain other domains.
- Reference other domains.
- Overlap by listing (managing) the same entities (devices).

- Share data (although, for security reasons, users may limit the use of data to a single domain).

EMA imposes no rules on a domain's composition or size. There are, however, practical limits on domain size relative to the amount of network management traffic a large domain can carry and the amount of storage it would require to support its Management Information Repository. Figures 2 and 3 provide examples of how management domains might be structured under EMA.

FROM DECnet TO EMA: THE EVOLUTION OF DIGITAL'S NETWORK MANAGEMENT

Since the advent of DECnet Phase III in 1980, Digital has provided embedded tools for monitoring and controlling DECnet operations. One example is the Network Control Program (NCP), a part of the VAX/

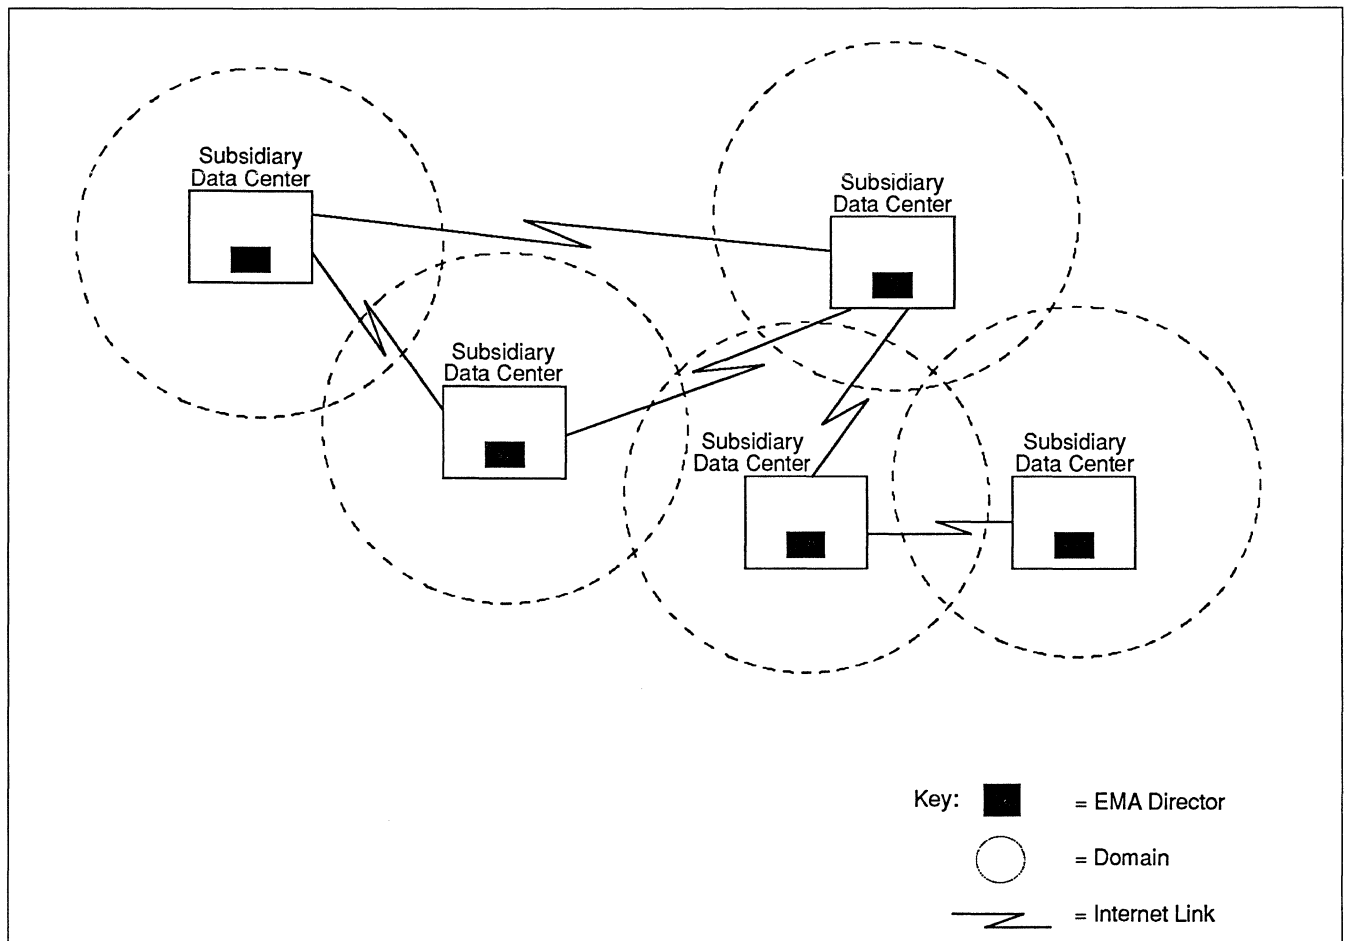


Figure 3. A distributed, multidomain configuration. Domains are defined according to corporate subsidiaries, located in scattered geographical locations. Each subsidiary network is an EMA domain and is managed as such. However, all of the subsidiaries share common network elements such as file transfer and electronic mail applications, as well as access to common internetwork communications facilities. Thus, it is useful to allow domains to exchange management information about the shared applications and communications facilities.

Digital Equipment Corporation Enterprise Management Architecture (EMA)

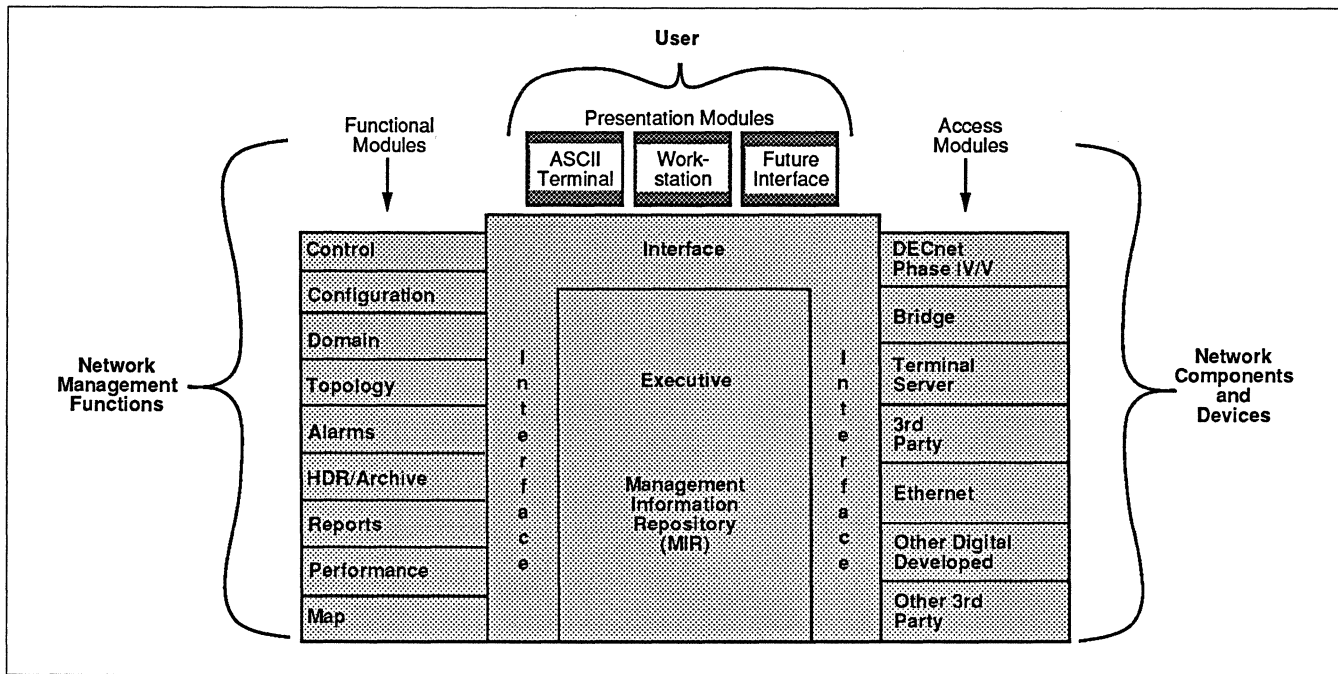


Figure 4. An extended illustration of EMA's structure, with examples of Function, Presentation, and Access Modules. Digital will incorporate existing DECnet product capabilities into the first release of EMA software products.

VMS operating system which manages DECnet-VAX, the Packetnet System Interface (P.S.I.) products, and other Digital software. While NCP and other embedded software tools were integrated with DECnet, they actually resided in discrete databases—collecting information and storing it locally in each DECnet node.

Some of the functions already built into DECnet included changing/examining network operation control parameters, generating test messages, examining and logging counters and events, and displaying/altering line or node status. DECnet provides adaptive routing and address changes, which makes it much easier to manage changing environments. Having these functions built in to DECnet is an advantage for Digital users, particularly since these management capabilities are well integrated with the VMS operating system. Digital has announced its intention to use these existing products to provide a migration path to future EMA products.

DECnet already possesses a layering scheme much like the OSI Management Framework. At OSI Layer 7, Digital defines a Network Management sublayer which uses the Network Information and Control Exchange (NICE) protocol for testing, downline loading, and upline dumping. Most of the network management functions performed by DECnet processors are concentrated in the Network Management Layer, although some remain in certain protocols such as the Digital Data Communications Message Protocol

(DDCMP). (For more information about DECnet's existing network management functions, see "Digital Equipment Corporation Digital Network Architecture (DNA) and DECnet," Report CMS60-384-101 in Datapro Reports on Communications Software.)

NICE, which is an openly published protocol, is functionally similar to OSI's CMIP. Digital will eventually replace NICE with CMIP in the next version of DECnet, DECnet/OSI Phase V.

In addition to embedded tools, Digital currently provides separated network management products including Remote Bridge Management Software (RBMS), Ethernim, Terminal Server Manager (TSM), Remote System Manager (RSM), the LAN Traffic Monitor (LTM), and the Network Management Control Center (NMCC)/DECnet Monitor. The DECnet Monitor is a set of sophisticated tools which respond to commands and present color graphics displays depicting the network's condition.

In spite of existing DECnet management capabilities and separate network management products, Digital still lacked some important features prior to the EMA announcement. First, Digital provided no integrated network management capability for logging and displaying events from other vendors' devices, particularly modem management systems. Second, no other major vendors supported Digital's proprietary DECnet management protocols (such as NICE).

Digital Equipment Corporation Enterprise Management Architecture (EMA)

EMA embodies Digital's strategy for providing integrated management—and for gathering third-party support. EMA is an ambitious undertaking for Digital; users should expect the process to take time in order for it to be done right. Indeed, Digital has

already delayed publication of the EMA interface specifications by several months. Those users who can afford to wait will be rewarded with a flexible, well-designed framework for managing the complexities of multivendor networks. (See Figure 4.) □

OpenView's Architectural Models

This report will help you to:

- Discover how OSI and de facto standard network management approaches can coexist.
 - Evaluate Hewlett-Packard's approach for migrating to OSI Management.
 - Apply organizational and operational models to evaluate other OSI-based network management products.
-
-

This report introduces two models developed as part of Hewlett-Packard's OpenView Network Management (NM) architecture: the Organizational Model and the Operational Model. Both models are potentially valuable tools in designing any NM product. They provide both a high-level view for initial planning, as well as a detailed view for implementation.

INTRODUCTION

Hewlett-Packard (HP) has developed two models as part of its OpenView Network Management (NM) architecture. The models have been found to be extremely useful tools to NM solution developers in the way they can capture a high-level view as well as a specific detailed view of the NM environment.

The *Organizational Model* is intended to assist designers in identifying management functions and their relationships to one another. In contrast, the *Operational Model* reveals sufficient design detail to support dataflow and coexistence analysis. Both models support multiple levels of integration, which allow

many more systems to be integrated under a common network management architecture with varying degrees of effort.

These models have been used to deal with distributed processes, multiple communication stacks, management supervision, and multiple user interfaces. Features of OpenView include multivendor management through standard(s), flexible physical realization, migration through coexistence, multilevel integration, and open access to network management services. This report introduces these models and some of the rationale behind their development.

Standards are the key to managing multivendor environments. HP has led the industry in its commitment to OSI standards. The problem HP (and others) face in fulfilling this commitment is how to introduce OSI based products into existing networks presently based on non-OSI protocols. OpenView addresses migration by supporting coexistence of both OSI and de facto standard NM solutions during indefinite periods of migration. The development of these models began with the ISO OSI Systems Management Architecture. The ISO work was studied carefully and pressed to its limits to see what (if anything) lay beyond its scope. The object oriented nature of the ISO work provided a sound foundation but some areas of a NM solution were not covered.

This Datapro report is based on "OpenView's Architectural Models," by Keith S. Klemba, Hewlett-Packard Information Networks Group, from the *IFIP Symposium on Integrated Network Management*, Boston MA, May 14-17, 1989. © May 1989 by Hewlett-Packard Company. Reprinted with permission.

OpenView's Architectural Models

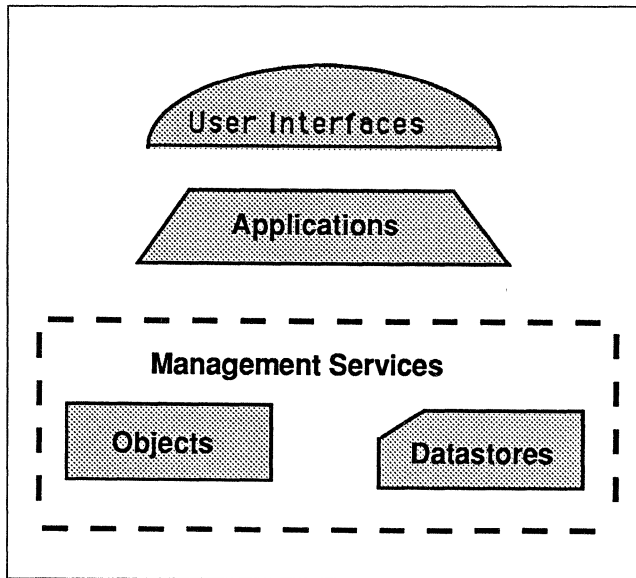


Figure 1. The OpenView Organizational Model.

ORGANIZATIONAL MODEL

The OpenView Organizational Model is made up of three major components: User Interfaces, Management Applications, and Management Services. Management Services are primarily provided by two subcomponents: Objects and Datastores. The Organizational Model uses these components to model the functional and relational composition of a NM solution.

These components are high-level and do not (in this model) represent the design-level or implementation-level scheme. During the analysis phase of developing a NM solution, it is useful to suppress these levels of detail (e.g., protocols, comm. profiles, and deployment). However, these are addressed later with the Operational Model.

Figure 1 shows the components of the Organizational Model. The shapes enable differentiation between components in actual solutions.

In developing the Organizational Model, both management and modeling technologies were used. For example, delegating authority among Management Services is a management technology, while use of object-oriented decomposition is an advanced modeling technology.

Perhaps the most important contribution of the Organizational Model is the building block concept for management. That is to say that NM solutions are designed such that managers are manageable. This is accomplished by putting the value-added portions of a NM solution into Management Services which can

then be linked into management chains. Consequently, we are using simply the term "Objects" rather than "Managed Objects" because it supports a broader use of the object-oriented decomposition.

Organizational Model Components

User Interface: Represents the technology used to connect the user with the NM solution.

Who are the users of NM? Where are they and how many of them are there? Are they all equal or can some do more than others? What types of skills or equipment must they have? These questions led to the creation of the User Interface component.

In this model, NM activities are conducted by Users accessing Management Applications via a User Interface. Examples of User Interfaces include devices such as terminals, PCs, Workstations, instrumentation panels, warning lamps, reset buttons, etc. Some of these devices (e.g., workstations), actually provide an entire environment for the user to access applications (e.g., a windows environment). In these cases, the User Interface component includes these environments.

Management Application: Represents the portion of a NM solution which supports a specific management activity through a specific User Interface.

A Management Application accepts user input and prepares information for display to the user. It makes use of available facilities to carry out its NM activity. These facilities include those offered by Management Services, User Interfaces, and native environment services (e.g., OS utilities). The details of communication services, normally prominent in a management Model, are suppressed in this model.

The Management Applications are under the direct control of the user and may be initiated or terminated by the user at any time. A given Management Application may support one or more User Interfaces.

In principle, Management Applications can use any of the Management Services necessary to carry out an activity, restricted only by security measures. Therefore, Management Services may be combined in limitless combinations to support new Management Applications. In practice, and especially during periods of migration, a Management Application can only use those services with which it is compatible.

Management Services: Represents NM support facilities provided chiefly by Objects and Datastores.

OpenView's Architectural Models

Management Services are the key components of the OpenView NM architecture. These services are designed to support Management Applications as well as other Management Services. The Management Services component is made up of Objects and Datastores.

There is an important distinction between Management Applications and the Management Services in this Model. Often a NM solution combines Applications and Services into a single module, locking key value-added information into the final display to the user. By separating this value-added information into Management Services (Objects) the information can be used by other NM solutions. Consequently, the OpenView NM architecture separates the Application function from the Management Service function in order to facilitate building upon NM solutions in an ever expanding open environment.

Objects: This subcomponent represents an object in the network that behaves in accordance with a registered specification. Anything which can be monitored and/or manipulated by modifying control algorithms and/or data through management protocols is considered an Object (Managed Object according to ISO). More precisely, it refers to the attributes, actions, and events described in an object oriented specification found in a registered Management Information Base (MIB). Objects include abstract Objects which may be created to represent a general class of Object (e.g., Modem).

Examples of Objects include computers, modems, gateways, bridges, X.25 switches, applications, LAN managers, subnet managers, global managers, etc.

Datastores. This subcomponent represents the data storage and retrieval services provided by management Datastores.

Organizational Model Example

Figure 2 is an example of applying the Organizational Model to a LAN bridge management solution, in which three different Management Applications make use of an infrastructure of Management Services.

Through an X-Window interface, the user accesses a network configuration application. This application makes use of the Management Services provided by a general configuration object and a configuration history datastore.

Through an OpenView-Windows interface, the user accesses both the network configuration and bridge exerciser applications. The bridge exerciser applica-

tion uses the Management Services provided by the bridge manager and bridge parameter object. It also uses the bridge history and bridge inventory datastore services.

Although somewhat limited, local buttons and LEDs permit the user to interact with a Management Application to view status and reset the bridge.

This solution could be developed over time. In early versions, the Management Services were used exclusively by a single Management Application. Gradually other NM solutions were developed that incorporated the Management Services of earlier solutions.

OPERATIONAL MODEL

In the previous section, OpenView's Organizational Model for management in a networking environment aided in identifying the functional elements and expressing them as components of a management solution. OpenView's Operational Model is a design-level model useful for illustrating how the components of the Organizational Model are deployed and how multiple NM solutions can coexist. It also provides a means for dataflow and management supervision analysis.

The ever-changing NM environment, brought on by changes in standards and technology, requires a high degree of flexibility in an Operational Model. The following discussion of OpenView's Operational Model provides descriptions of its components and how the model might be used.

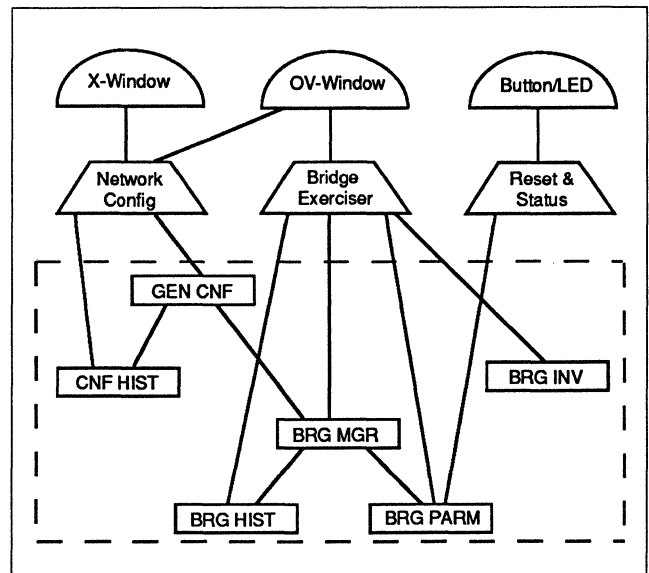


Figure 2. Organizational Model Example.

OpenView's Architectural Models

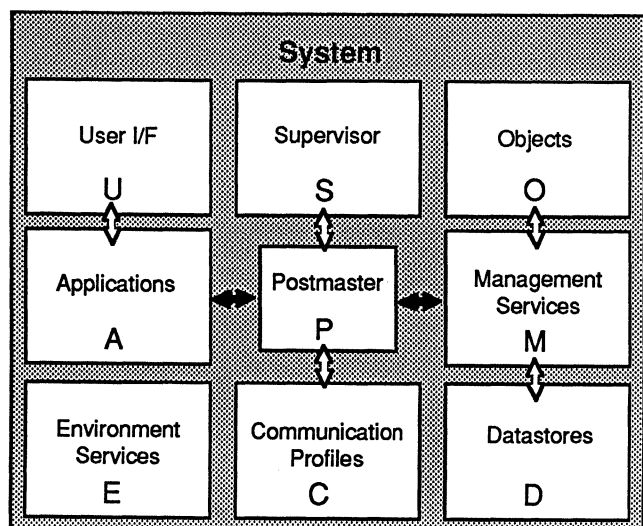


Figure 3. OpenView's Operational Model.

Operational Model Components

Figure 3 shows the components of the Operational, or "Nine-Squares" Model.

Outer Encompassing Box: The outer box represents a single physical system. The portion of the NM solution provided by this system is partitioned into the nine components located within this system.

Arrows: The solid arrows represent standardized interfaces within the architecture. The hollow arrows represent interfaces within the architecture which are system or implementation specific.

User Interface (U): Represents the man-machine interface facilities within a system. See Organizational Model.

Management Application (A): Represents the portion of a solution which supports a specific management activity through a specific User Interface. See Organizational Model.

Environment Services (E): Represents the general facilities within a systems environment which may be used directly by any of the other components. Examples of this component include system utilities such as File Sorting, Terminal Emulation, and File Transfers.

Supervisor (S): Represents the management Supervision present within a system, which is responsible for maintaining the components in this system. Examples of the Supervisor include the maintenance of Management Services software within this system and their infrastructure.

Postmaster (P): Represents a basic object-oriented message routing facility. It operates as a message switch, determining its routing action from an object-oriented routing table. The routing table is under the control of the Supervisor of the system. It does not impose any modifications to existing protocols. Great care must be taken NOT to add additional services to (P), tempted by its central influence.

Communication Profile (C): Represents a set of communication services organized to support the transaction-based message routing of (P). Each (C) is a complete profile or "path" for communicating NM information with other systems. The profile selection can be done explicitly or implicitly when the message is given to (P). The routing table (P) used to support this function may be influenced by (S) to redirect messages or override explicit choices.

Each system chooses which profiles it will support and how (e.g., dynamically, statically, multiplexed, individual). (M), (S), (D), or (A) components may extend an existing profile by adding levels of communication services.

Objects (O): Represents the Objects located in this system. See Organizational Model.

Management Services (M): Represents all Management Services available in this system. It assumes the Management Services of all Objects and Datastores within this system. Consequently, it is not necessary to repeat the name of an Object or Datastore in the (M) box which appears in the (O) or (D) box. The (M) box can be used to name management services within this system which are neither Objects or Datastores yet have been made a part of a NM solution and conform to the architected (P)/(M) interface. See Organizational Model.

Datastore (D): Represents management data storage accessible in this system. Its presence in the model simply gives evidence of its existence and information content. The details of its implementation are outside the scope of this model. See Organizational Model.

Operational Model Example

Figure 4 shows the OpenView Operational Model being applied to the same LAN bridge management solution used in Figure 2. Notice the ease with which the Organizational Model translates into the Operational Model.

The Bridge itself is modeled as a system. It is equipped with a reset button and status LEDs which

OpenView's Architectural Models

constitute a fundamental User Interface (U) for managing this bridge. The software (and/or board logic) which controls the LEDs and responds to the reset button is modeled as a Management Application (A). The bridge is also equipped with some ability to receive control messages from the LAN. The communication protocol for exchange of these messages is modeled as a Communication Profile (C). The specific parameters of the bridge which can be controlled or interrogated appear as a Bridge-Parameters Object (O). The actual message processing (e.g., Get or Set Bridge-Parameters) is modeled as Management Service (M) provided by the Bridge-Parameters Object. Because this bridge implementation does not make use of the Postmaster Component, (P) is omitted, likewise are (S), (D), and (E).

The middle system in Figure 4 is a PC based management station on the LAN which is capable of managing this bridge by means of control messages. OpenView Windows provides the User Interface (U) on this system. Imagine that the PC management station was developed to provide configuration management for all the bridges on the LAN. To accomplish this function, a Bridge-Manager Object (O) is created on the PC and given a set of services (M) such as "Set or Set Config Data". A single Management Application (A), is developed, allowing the user to execute these services and view their results.

In order to support communications between the Bridge-Manager Object and the Bridge-Parameters Object, the PC must provide a common Communication Profile (C) between the bridge and the PC.

The PC implementation does support a Postmaster (P) for object-oriented message handling. Bridge-History and Bridge-Configuration Databases (D) are resident on the PC's hard disk. The services provided

by these Databases are modeled as (M). A Supervisor (S) has been provided to allow for interrogation and control of the Postmaster's routing tables. At least one Environment Service (E), an X-Windows Client, is supported on the PC.

The third system in Figure 4 is a workstation using X-Windows as the User Interface (U). A General-Configuration Object (O) is created on the workstation, with services (M) that provide the means to identify the specific Bridge-Manager Object responsible for a given Bridge-Parameters Object. One of the two workstation Management Applications (A) allows the user to exercise the services of the General-Configuration Object. The other allows the user to exercise the services of the Bridge-Manager Object. Just as in the PC, the workstation implementation makes use of a Postmaster (P). There are likely to be many more components in the workstation (e.g., [E], [M], [S], [C], and [D]). This example has highlighted only a few.

The Management Services (M) provided on the PC management platform are also used by the workstation. In order to accomplish PC and workstation integration, a second Communications Profile (C) is added to the PC, which supports ISO/OSI Common Management Information Protocol (CMIP). An identical Communication Profile (C) is added to the workstation; it is likely that the workstation already supports CMIP.

To illustrate the flexibility of this architecture, imagine that the Bridge-Manager Object (O & M) is moved off the PC and on to the workstation. Such a move might be necessary if, for example, the PC is not dedicated to its bridge management role, thus isolating the workstation from the bridge manager object. The Postmaster on the PC would be advised

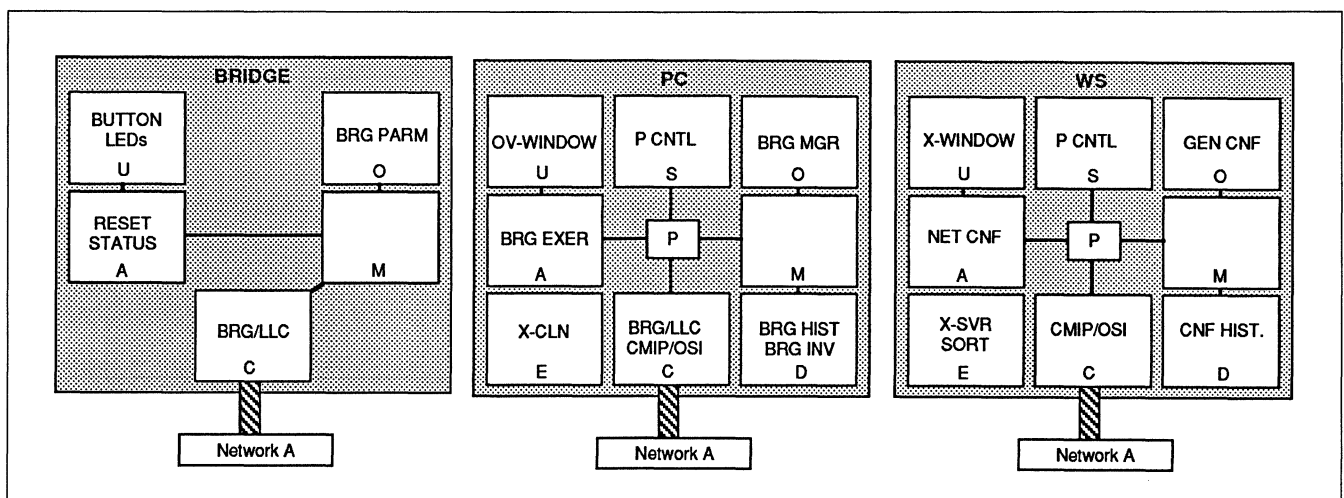


Figure 4. Operational Model Example.

OpenView's Architectural Models

of this move and the request for these services coming from the (A) on the PC will be routed to the (M) on the workstation. An alternate Communication Profile (C) may also be added to both the PC and the workstation to strengthen their integration. In spite of these changes, the PC user is still able to access and run the Bridge-Exerciser Application on the PC.

A new NM solution could be built on this management infrastructure making use of any of these Management Services. Furthermore, the deployment and administration of the new solution can also build upon these systems.

CONCLUSION

Degrees of freedom in NM solution development are useful to product planning, design, and implementation teams. Key to achieving this flexibility is the decomposition of a NM solution into modeled components, and the adoption of an integrating component such as the Postmaster (P) in OpenView's Operational Model. OSI NM might be possible without these, however, the OpenView architecture suggests a strategy for development which provides leveraging and coexistence of both OSI and non-OSI management solutions.

This report has shown how the OpenView Organizational and Operational Models can move NM solutions from analysis to design. Key to the success of these models has been the decomposition of a NM solution into component parts, the most important being that of Management Services. By putting the

value-added service of a NM solution into Management Services provided by Objects, future NM solutions will be able to build upon these services. Furthermore, by insulating Management Services from the Communications Profiles and User Interfaces, coexistence and migration do not disrupt acquired Management Services and dependencies.

Having gained momentum on synthesizing the OpenView management Architecture, HP recognizes the need for an Information Model which addresses the coordination of NM information being brought together from diverse NM solutions through combined Management Services. It also addresses the Datastore requirements necessary to support the object management infrastructure. The Information Model of OpenView is the focus of current development work at Hewlett-Packard.

REFERENCES

- Information Processing Systems—OSI—Management Information Protocol Specification—Part 2: Common Management Information Protocol, ISO/IEC DIS 9596-2, Oct 4, 1988.
- Information Processing Systems—OSI—Basic Reference Model—Part 4: Management Framework, ISO/IEC DIS 7498-4, Jan 21, 1988.
- International Telegraph and Telephone Consultative Committee (CCITT), IXth Plenary Assembly—Document 31—Study Group IV—Report R26—New Recommendations of the "M" Series, AP IX-31-E, Apr 1988.
- Advanced Networked Systems Architecture (ANSA) Reference Manual, Rel.00.03, Jun 1987. □

IBM SNA and NetView

This report will help you to:

- Grasp SNA network management fundamentals.
 - Compare IBM's network management approach to those of its competitors.
 - Use NetView's strengths to your best advantage and select alternatives to compensate for its weaknesses.
-
-

NetView holds the undisputed top spot in the market for integrated SNA network management packages. NetView's most direct competitor is Cincom System's Net/Master, which provides similar SNA network management functionality with some added benefits. NetView now holds a commanding lead over Net/Master, however, and NetView installations outnumber Net/Master's by at least 3 to 1. Within the past 12 months, both AT&T and Digital Equipment Corporation have announced major network management platforms that threaten to reduce IBM's commanding lead in the market.

How does NetView manage SNA networks, and what are its strengths and weaknesses? How will NetView address the growing demand for standards-based connectivity? This report provides answers to these questions by examining IBM's approach to network management within the context of SNA and the emerging OSI standards.

IBM'S APPROACH TO NETWORK MANAGEMENT

IBM's network management strategy emphasizes four points:

- **Centralization**—This is appropriate for managing the hierarchical structure imposed by SNA.

- **Data Networks**—Although the recently announced NetView Call Accounting and Voice Network Definer products call attention to voice, NetView's real strength (both in terms of technology and market share) lies in data networks.

- **The Logical Network**—The logical network is represented by the *applications traffic* passing over the physical links. NetView provides sophisticated logical network management; its capability to manage the physical network (devices, circuits, etc.) is far less comprehensive.

- **Integration**—NetView provides a common operations interface to manage local and remote components (both logical and physical) of large corporate networks.

The following sections describe the basis for these four points.

Centralization

Most communications network architectures distribute overhead away from the mainframe. In particular, Digital's Digital Network Architecture (DNA) emphasizes remote points of control (see "Digital Equipment Corporation Enterprise Management Architecture (EMA)," Report NM40-325-101).

IBM SNA and NetView

SNA, however, uses mainframe-based software to control virtually the entire communications path and the associated processing. Host-based control extends in both directions, from the input/output statements of an applications program to the printer or screen of a user's terminal. During the past three years, however, IBM has modified its rigidly hierarchical networks by introducing Advanced Peer-to-Peer Networking (APPN) and other subarchitectures. (For more information, see "Network Management in APPN Networks," Report NM50-100-301.)

Despite these modifications, SNA still requires the creation and maintenance of extensive network configuration parameters within various mainframe software components. Much of this configuration information must be maintained by systems programmers.

Most SNA processing occurs within the telecommunications access methods (typically VTAM) which reside on IBM mainframes. Therefore, the software that manages SNA networks (NetView) resides on the mainframe as well. A complete description of NetView's centralized, hierarchical structure is provided in this report within the section entitled "NetView's Structure."

Data Networks

IBM has dominated the data processing industry for over 25 years. IBM knows how to develop data products and, more importantly, sell them. IBM is unbeatable when it comes to maintaining account control among MIS managers; but IBM's image among telecommunications managers leaves a lot of room for improvement. (See "IBM's Network Management Strategy," Report NM60-491-101.)

IBM introduced three voice network management products when it unveiled NetView in 1986. While NetView has emerged as a major de facto standard, the voice products were never successful in terms of sales or price/performance.

With its new ACCUMASTER Integrator, AT&T is now in a position to capitalize on IBM's weakness in voice network management. (For more information, see "AT&T's Unified Network Management Architecture," Report NM40-313-101.)

The Logical Network

SNA outlines the logical structure of an IBM network, the configuration of its nodes and the links between them, and their relationship in the overall

hierarchy of control. SNA also outlines the logical process of communications through a layered architecture of processes, through which a message—the fundamental, generic unit of transmitted information—must pass on its way from source to destination.

All communications in an SNA network occur within *sessions*. A session is a logical, two-way connection between two network addressable units (NAUs). An NAU is a segment of system software that represents a specific device or application program to the network.

NetView provides comprehensive monitoring and control of logical sessions. Other vendors' products outperform NetView in physical network management, particularly transport management. Transport management involves monitoring and controlling circuits and related facilities. (For more information, see "Transport Management," Report NM20-400-201.)

Integration

Prior to NetView's release in 1986, no major network management products provided a common user interface for the proliferation of separate network management products that typically reside on large corporate networks. NetView, along with Cincom's U.S. release of Net/Master around that same time, introduced integrated network management capabilities to the market.

NetView provides a degree of *integrated* network management because it allows the user to monitor both physical and logical aspects of the network from one console. (AT&T's new ACCUMASTER Integrator provides the same type of integration.) As yet, neither NetView nor any other products comprehensively integrate multivendor device alerts with inventory, configuration, and historical information (stored in a central repository). Ideally, this type of integration would permit comprehensive automatic filtering of unnecessary messages and low-level processing of certain events.

While its integration capabilities are far from ideal, NetView does provide sophisticated monitoring and control of the logical network and some degree of physical monitoring as well (status of devices, circuits, etc.). These capabilities are brought together at one centralized NetView console.

IBM SNA and NetView

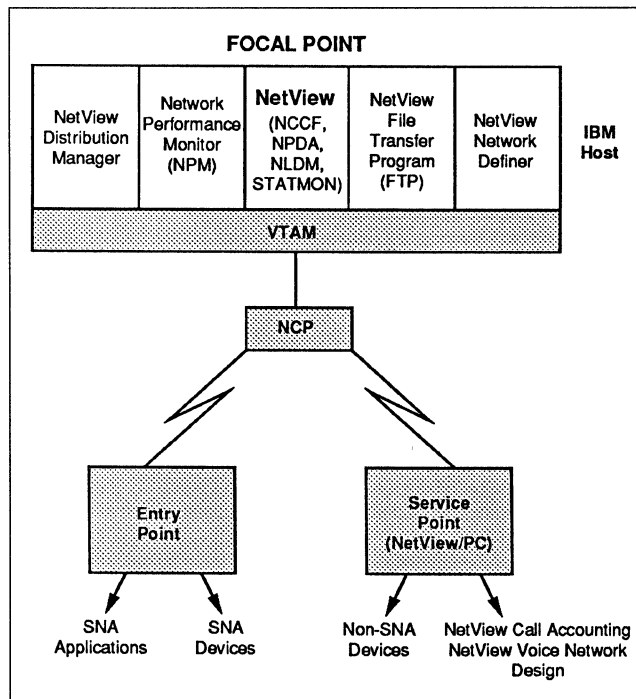


Figure 1. This figure illustrates NetView's hierarchical three-point architecture. Centralized monitoring and control is provided by focal point products, which reside in the host. Entry points forward alerts from SNA network components to the host. Service points, a type of entry point, forward alerts from non-SNA devices to the host.

NETVIEW'S STRUCTURE

NetView embodies a hierarchical three-point architecture (see Figure 1). At the top of the hierarchy, the *focal point* provides the functions necessary for managing the network (both local and remote components) from a central location. Focal point products, such as NetView (NCCF, NPDA, NDLM) and NPM reside in the mainframe host.

Entry points collect management data (alerts, etc.) from distributed SNA nodes and forward it to the focal point for centralized processing. Entry points manage attached SNA devices by executing commands sent from the focal point.

A *service point* is a special type of entry point which forwards alerts from non-SNA devices (i.e., devices not supported by entry points.) NetView/PC, IBM's primary service point program, uses Network Management Vector Transport (NMVT) formats, a type of SNA alert. NMVTs specify alert origin, severity, probable cause, and recommended actions.

NETVIEW COMPONENTS

NetView actually comprises a set of host-resident programs for network operation, error detection, error correction, and management. As discussed previously, IBM's approach dictates centrally controlled configuration and error reporting for every aspect of the network. This includes both physical (devices) and logical (session), from the logical unit (LU) at one end of a session (the host-resident application) to the LU at the other end (the display terminal screen). To this end, all SNA devices, including IBM modems, contain some facility for detecting and reporting errors.

NetView: Three Old Products under One New Umbrella

With the introduction of NetView Release 1 in May 1986, IBM combined into a single product all of the functions of three existing separate products:

- Network Communications Control Facility (NCCF)
- Network Problem Determination Application (NPDA)
- Network Logical Data Management (NLDM)

NCCF, now called the NetView Command Facility, runs as an applications subsystem—much like CICS—with ACF/VTAM or ACF/TCAM. The Command Facility includes the operator interface and network logging facilities. It allows operators to enter any NetView command, Command List (CLIST), or VTAM command from any designated NetView terminal (IBM 3270 or compatibles). In this manner, operators can issue configuration and diagnostic commands to SNA network devices.

NCCF also serves as a base for IBM's two network troubleshooting programs, NLDM and NPDA.

NLDM, now called the NetView Session Monitor, collects performance statistics about the sessions and routes defined in SNA's upper layers and helps to identify network programs. The Session Monitor manages the logical network, consisting of traffic patterns (rather than physical circuits.)

The Session Monitor runs as an application under The Command Facility. It records failures and degrading conditions on all logical SNA sessions active on the network at any given time. In operation, The Command Facility records certain relevant data on every session: the logical units participating, the session type, and the class of service. When it

IBM SNA and NetView

detects problems or operator-issued commands, it records all header and trailer information from the specified session's message units.

The Command Facility is most useful in tracing the causes of lost data and in resolving protocol incompatibilities between communicating LUs. By collecting relevant resource activity data during and just prior to failure, the Session Monitor assists in network problem determination. All failure information shows probable causes and recommended actions.

The Command Facility displays data in full-screen mode using seven colors to differentiate various resources.

NPDA, now called the Hardware Monitor, also runs as an application under The Command Facility. It records failures and degrading conditions on the physical SNA network and can initiate trace programs to find the sources of SNA network hardware problems. The Hardware Monitor operates through SSCP-PU sessions between the host access method and the network's Physical Units (PUs).

The Hardware Monitor provides a series of panels to help the operator detect, diagnose, and resolve problems in SNA's two lower layers, which consist of the physical interfaces (modems, cluster controllers, terminals) and data routes (lines, etc.) that link local and remote devices.

NetView Add-Ons

In addition to NCCF, NLDA, and NPDM, NetView also provides some of the functions of the VTAM Node Control Application (VNCA) and Network Management Productivity Facility (NMPF). NMPF is a set of job streams, programs, and data sets that assist operations staff in installing, learning, and using other NetView components.

NetView represents more than a repackaging of existing products, however. NetView gives IBM a logical framework for enhancing current products and developing future tools. IBM also effectively uses the term "NetView" in its marketing strategy to promote separate systems management and networking tools under the NetView umbrella, some of which are discussed below.

Network Performance Monitor (NPM). NPM is a separate, host-resident program that replaced the Network Performance Analyzer (NPA) and VTAMPARS II. NPM gathers statistics on line usage, message vol-

ume, and response times. It forwards this data to the mainframe for use in preparing realtime reports.

NPM is a VTAM system application that can be invoked from a NetView operator terminal console via the Terminal Access Facility (TAF).

Status Monitor (STATMON). STATMON, based on the same concepts found in VNCA, is a front-end interface that displays the status of the network resources hierarchically and allows on-line viewing of the network log. A full-screen, color-coded panel displays the status of a domain under the control of the network operator. The operator can view all or selected Network Control Programs (NCPs) in the domain and proceed down to all lines in a specific NCP, to all Physical Units (PUs) on a specific line on the NCP, and so forth. The color code of a resource indicates its status as active, disabled by the operator, never active, or malfunctioning.

ISCF extends the system operations support of NetView to control and monitor multiple automation target systems from a centralized controlling system. It runs as a NetView application and operates in conjunction with the ISCF/PC program offering. This program provides users with access to NetView's operational facilities, enabling them to automate target system operations such as system initialization procedures.

NetView/Access is an MVS licensed program that enables SNA network users connected to SNA Applications Monitor (SAMON) to use a number of applications concurrently from a single terminal. NetView/Access protects the network and the applications against unauthorized use and performs automatic logon to, and logoff from, those applications the user is authorized to access based on existing user profiles.

SAMON, along with NetView/Access, extends the NetView access services capabilities. It enables terminals to display the status of all VTAM applications within a network and to connect the terminal to one of the applications. In addition, it provides status information about both the individual application and the entire network.

Network Definer is an interactive application designed to reduce the effort and skills required to create and update definition tables for VM-based SNA networks, especially in the 9370 environment. The Network Definer provides full-screen, menu-oriented dialogs for network definition. It runs under VM/SP and will create network definitions for ACF/VTAM, NetView, and Remote Spooling Communications Subsystem (RSCS) Networking using IBM-supplied

IBM SNA and NetView

defaults or user-supplied input. The user can create network definitions at a central site for later distribution to remote systems. Network Definer will automatically distribute subsequent updates of these definitions if an active RSCS connection to the target remote systems is available.

This NetView facility allows a central site to perform timed updates and create end-to-end routines so that the entire network is fully connected. A sample program is provided and can be tailored to meet local needs. Entire network descriptions can be copied using the supplied dialogs.

NetView File Transfer Program provides high-performance bulk data transfer between SNA/370 systems. It is structured into a Base product and an Advanced Function Feature. FTP is an interactive VTAM application that can transfer bulk files between any MVS-, VSE-, or VM/SP-based installation in an SNA network. File-to-file transfer can be accomplished without requiring spooling.

NetView Distribution Manager provides services for centrally controlled data distribution and the implementation of software changes in SNA networks composed of a variety of distributed/departmental systems. The Distribution Manager supports SNA networks composed of VM end nodes; System/36 end nodes connected through a System/36 intermediate node; PCs and PS/2 end nodes connected through a System/36 intermediate node; and System/36, VSE, 4680, Series/1, Series/1-PC Connect, or 8100 directly connected end nodes.

NetView Call Accounting and NetView Voice Network Design. These voice products, announced in the spring of 1989, replace existing NetView programs that failed to capture any sizable customer base. Both products were designed by third-party vendors, not IBM.

Both NetView Call Accounting (developed by DMW Communications, Inc.) and NetView Voice Network Design (developed by Vector Software, Inc.) work with NetView/PC. Although the new products carry the NetView name, neither product is actually integrated with NetView (once again, IBM uses the term "NetView" to enhance its products' marketability).

NetView/PC

Introduced shortly after NetView, NetView/PC is a multitasking personal computer subsystem which extends passive NetView *monitoring* functions to IBM

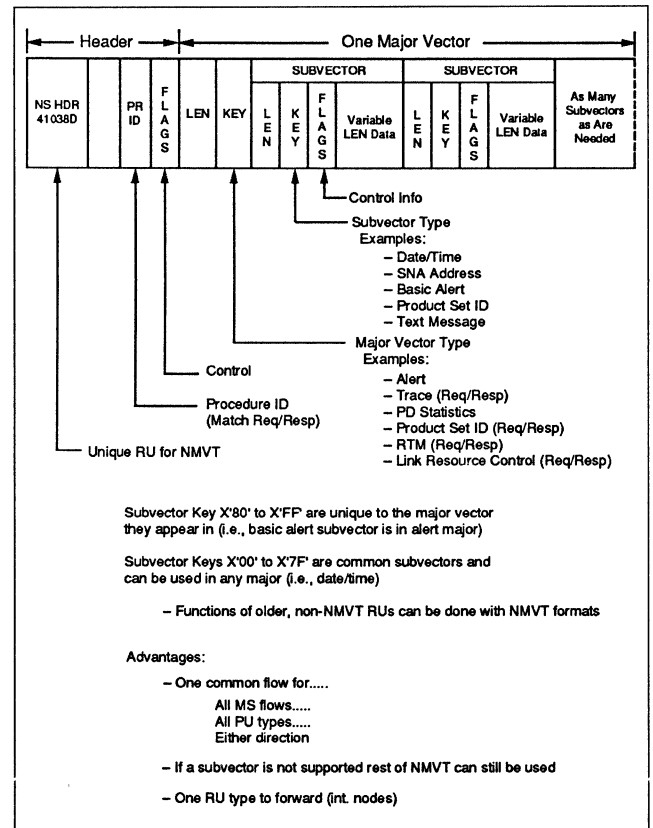


Figure 2. The Network Management Vector Table (NMVT) will become the key management services request format within SNA—used to support both solicited and unsolicited requests and replies between the System Services Control Point (SSCP) and the Physical Units (PUs).

Token-Ring Networks, IBM/Rolm CBXs, and other communications systems or programs developed by third-party vendors.

NetView/PC provides three types of communications using an SNA host. The first type provides applications programs with the means to forward generic alerts that do not require unique support in NetView. The application creating the alert provides the routing instructions: NetView/PC then routes the alert to the NetView/PC operator or to NetView. This alert structure follows the Network Management Vector Transport (NMVT) format and flows over an SSCP-PU 2.0 session path to NetView. (See Figure 2.)

The second type of host communications is the transfer of data files between the host and NetView/PC, using a Logical Unit (LU) 6.2 session. This session is supported at the host and by a CICS/Distributed Data Management (DDM) transaction program. Files can be transferred in either direction, initiated by the NetView/PC operator or by any NetView/PC application.

IBM SNA and NetView

Finally, the Service Point Command Facility allows a NetView-issued command to be routed to an application running in NetView/PC, where the command will be executed. The functions supported by the application, via this facility, are determined by the application developer and the capabilities of the application-managed device.

IBM provides NetView/PC interface specifications to interested parties. While IBM claims that over 40 vendors plan to develop NetView/PC interfaces, only a handful of products are on the market today. This lack of support is chiefly because NetView/PC is awkward and expensive to implement and, in effect, treats the third-party vendor's equipment as a nonintelligent device—even if the equipment possesses local network management capabilities.

NetView/PC Release 1.1 is currently available; Release 1.2 is due out later this year. According to IBM, Release 1.2 will run under OS/2 Extended Edition instead of DOS. OS/2 EE provides better support for multitasking—and many analysts believe that this will solve a number of NetView/PC's current problems. For example, the system will be equipped to handle more alarms; the product's current alarm processing capability has been a limitation in the past.

NetView Release 3

In September 1988, IBM announced NetView Release 3. This product will become generally available in 1990. The major enhancements include:

- **REXX Support.** Programmers will be able to write CLISTs in REXX instead of Assembler to customize NetView for their network's particular needs. REXX is a high-level string-processing language developed by IBM to support VM systems programming. While REXX is certainly easier to use than Assembler, industry experts question the appropriateness of selecting a string-processing language that few programmers know and tacking it onto NetView.
- **Knowledge Tool Support.** Knowledge Tool is IBM's expert systems product. Incorporation of Knowledge Tool will assist users in filtering messages and, thereby (they hope) in reducing message traffic.
- **Eventual OSI (CMIP) Support.** Although OSI support represents a major change in direction for IBM, it's not too surprising considering Digital's EMA announcement and AT&T's pending ACCUMASTER unveiling. The real dilemma is how IBM

will implement OSI support. (See the section of this report entitled "OSI and Future Directions" for more information.)

- **IBM SolutionPac NetView.** This offering, subtitled "Automated Network Operations," provides tailored network operation screens and command lists to help automate network operations. SolutionPac is actually a combination of services and software. The services include implementation planning as well as installation, tailoring, and testing of NetView. On-site training sessions are also included. An Automated Network Resource Manager software product, available only with SolutionPac NetView, identifies failing SNA network resources and automatically attempts to recover them. Two additional software products, the Dynamic Variable Generator and the Automated Help Desk, are also included with SolutionPac NetView.

Noticeably absent from the NetView Release 3 announcement was the mention of any central repository for network management data.

NETVIEW'S STRENGTHS AND WEAKNESSES

NetView's greatest advantage is simple—it's IBM's strategic product for network management. IBM dominates the mainframe and PC markets and is gaining market share in the departmental systems market with its AS/400. Furthermore, tying these three environments together (both technologically and from a selling standpoint) through its Systems Application Architecture (SAA) provides IBM with a legitimate claim on the direction of future developments in network control.

NetView provides sophisticated SNA network management capabilities. This, in itself, is a big money-maker, as there are reportedly 30,000 SNA licenses worldwide. NetView's strength in managing SNA networks is matched by its weakness in managing non-SNA data networks. Few vendors are actually marketing NetView/PC interfaces for their products, and not many users are buying them. NetView also lacks sophisticated voice network management capabilities, including monitoring and control of transmission facilities (public or private). NetView does not provide adequate physical network management of hybrid and miscellaneous network elements (non-SNA devices.)

IBM SNA and NetView

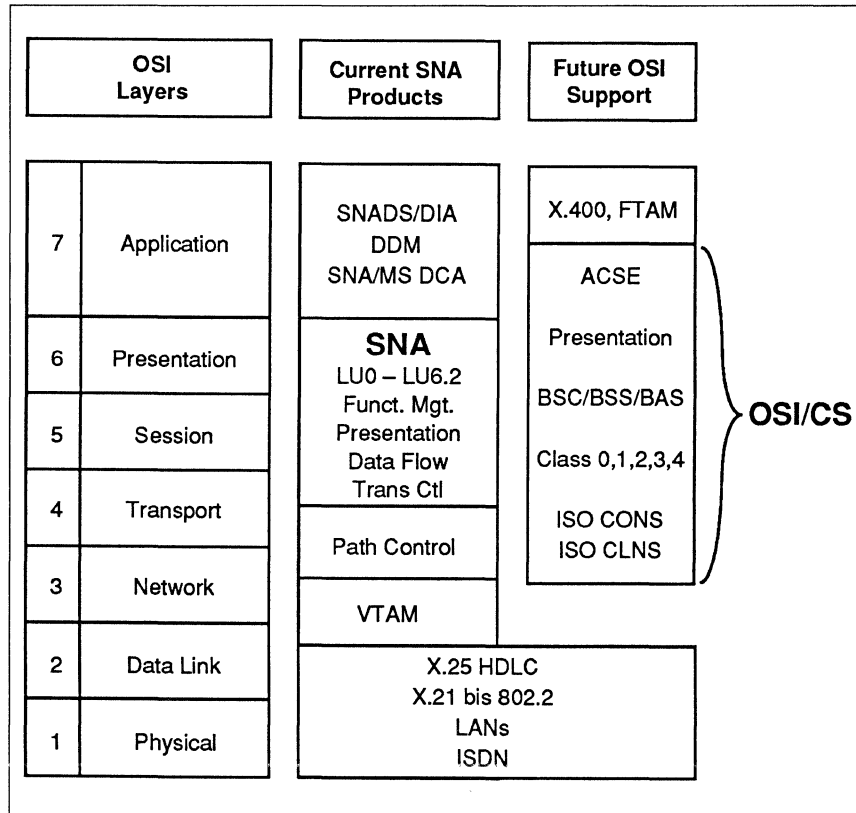


Figure 3. In September 1988, IBM announced support for selected OSI protocols and services. IBM intends to implement much of this support in the OSI Communications Subsystem, which will be available for MVS systems in March 1990. This diagram shows the relationship of the OSI/CS to other SNA layers.

OSI AND FUTURE DIRECTIONS

NetView provides a clear, organized migration path from the many different SNA control functions and programs. At the same time, Netview offers an open architecture with links to competitive hardware and software control elements. However, this is still a long way from integrated network management or an integrated network information database, especially for management elements from different vendors. Users and vendors alike are watching IBM to see how it evolves NetView in light of OSI's network management model.

Users (the U.S. government, in particular) are increasingly looking to OSI Management to provide guidelines for managing the complexities of large, multivendor networks. While more SNA networks exist than any other single type, few of these networks are "purely" SNA.

In September 1988, IBM announced that it would support an OSI Communications Subsystem (CS) in future strategic products, including NetView. IBM has pledged to perform a difficult task, since OSI and SNA were designed with different objectives.

SNA is a layered architecture that is similar, but not identical with the ISO's OSI model. Specifications

that correspond to OSI's Application Layer are now imbedded in IBM's SAA, but not in SNA proper. SNA's seven layers correspond roughly to the seven OSI layers; however, there are differences in handling routing and in defining the boundary between the lower "transport system" layers and the upper "logical control" layers.

IBM's OSI Communication Subsystem resides approximately within layers 3, 4, 5, and 6 (and part of 7) in the OSI model. (See Figure 3.) OSI/CS is a program product for MVS and VS, available in mid-1990. OSI/CS is a VTAM application which supports OSI protocols listed in Figure 2.

OSI/CS will serve as an OSI service point within the NetView architecture. Thus, NetView will support OSI Management in a master/slave relationship (defeating the purpose of the peer-to-peer model). OSI devices and systems will be under the centralized control of NetView. If implemented like the NetView/PC service point, OSI/CS will initially only transmit alerts from OSI devices; NetView will not be capable to receiving nor processing commands from OSI systems with network management capabilities.

Given the knowledge available to date, NetView's OSI support seems difficult and expensive to implement and not really worth the trouble. For IBM,

IBM SNA and NetView

however, OSI support will probably be enough to convince the U.S. government that NetView can, if necessary, support OSI Management and thereby comply with the government's OSI (GOSIP) mandate.

Although IBM appears willing to enhance its products to support industry-wide communications standards,

it is unlikely to redevelop products from scratch to make that support elegant or inexpensive. Nor does IBM need to do so at this time. SNA remains the dominant de facto standard for computer communications for the foreseeable future. Consequently, NetView (even in its present form) will continue to hold on to its lead in the market for SNA network management products. □

Network Management and Control Systems

This report will help you to:

- Investigate telco-provided network management and control services.
 - Evaluate a NYNEX model for integrating and optimizing multiple subnet control systems.
-
-

The concept of network management and control has evolved from a set of tasks that relate to the monitoring and management of the existing twisted pair network to a set of more complex tools that drive the performance of the data networks of today. The primary driver in the new method of attacking the network management problem is the growth of computer communications and its divergence from just a simple hierarchical network to a much more complex networking concept.

The networking of large-scale computer networks requires not only an understanding of the local and long distance telco plant, but also an understanding of the need of the total communications network. These needs are much more complex than those of the voice user of the system. Significant differences are those in response time and the ability to perform complicated networking. In addition, the need for higher reliability circuits, combined with the pressure of reduced costs, lead to the trade-off of having alternate routing and automatic restoral procedures available.

Many companies have provided communications systems that may at times be more expensive than the alternative telephone-based system. However, these systems often do not include a sophisticated Network Management and Control System (NMCS) capability such as that provided by NYNEX that the end users value more than the communications. Thus, the asset

of such a company is not in the telecommunications network alone, but in the NMCS capability and resources to handle and support the end user requirements and needs.

This report focuses on the concept of providing the data and computer communications user with a set of facilities that allow for the integration of many of the disparate tools that are presently available into one single network management utility.

In particular, in this report we develop a new model for NMCSs that expands the capability to support a sophisticated end user in a multivendor environment and allows for the integration and optimization of multiple subnet control systems.

PROBLEM DEFINITION

The basic problem that is addressable by an effective network management and control system is easily stated. Simply put, an effective NMCS allows the end user to have a single point of contact for problem resolution whenever the user hits a key on the terminal and the response is not as expected. Such a definition is all encompassing and it is not anticipated that a solution to this problem is readily forthcoming. However, it is essential to understand the ultimate goal of network management and control.

This Datapro report is based on "Network Management and Control Systems," by Terrence P. McGarty and Larry L. Ball, NYNEX, from the IEEE 1988 Network Operations and Management Symposium, New Orleans, LA, February 28-March 2, 1988. © IEEE Communications Society. Reprinted with permission.

Network Management and Control Systems

Computer Communications

In the world of computer communications, the user sees a different set of problems than that of the voice user. The data manager is faced with the overall management of the computer users' needs to ensure total end-to-end integrity of the process. This integrity is less tolerant of the errors that occur in the communications channel and also reflects the delays in the processing and transport of the data signals.

For example, in an IBM SNA environment, the end user may have a network that is composed of the following elements:

- Host IBM Model 3091
- Front End Processor: 3725
- Modems: Codex
- Local leased RBOC line at 56 Kbps using X.25
- InterLATA Network at T1 rates using a Cohesive controller
- Local RBOC multidrop line
- Modem: Codex
- Cluster Controller: 3274

End User Requirements

The end user requirements are generally similar in a computer communications environment. Consider a user in a large IBM-based environment. The user establishes a session under an SNA environment and the connection is made between the 3725 front end processor and the 3274 cluster controller. The lines may be a set of polled multidrop lines operating at 9600 bits per second. The end user utilizes a standard 3270 display terminal.

In this environment, the end user expects the following:

- A rapid response time to all keyed requests. The response time for the last character in to the first character out should be less than 2.5 seconds 99% of the time.
- An error rate on the line that ensures overall system performance. The line error rate typically should be less than 1 in 10^6 .

- In the event of a line or system trouble, the user should have access to an 800 type number system that allows for rapid problem determination and restoral. In particular, no problem should last longer than 10 minutes.
- The problem reporting and the resolution should appear as a seamless process to the end user. The end user should not be made aware of the sets of players that may be in between the overall network of users.

Performance Requirements

The performance requirements for an NMCS system are based on the needs of the end user in terms of network availability and time to restoral. The specific requirements may be modified when there are needs of the underlying network that result in increased levels of errors and outages.

Typical performance requirements are as follows:

- Percent of undetected user faults. These are the faults that are detected by the user rather than by the system. Thus, if 5 percent of the total faults are those that were first detected by the user rather than by the system, this may represent an unacceptable level of faults.
- Response time to fault isolation. This represents the time it takes to isolate a fault once it is reported. The system should have the capability to isolate faults to the source or cause in typically 5 minutes or less.
- Response time to resource reallocation. This is the time it takes for the network to respond with backup resources to allow for continued operation. This may vary dramatically from user to user. Some users may be willing to pay for total redundancy in the network and this time may be a fraction of a second. On the other hand users may not have direct redundancy in the network and alternate routing may be necessary.
- Response time to user complaint. The system must be able to respond to a user query in less than several seconds. This includes only the answering of the phone for a typical request. In addition, the system must be able to provide the customer service representative with the direct access to the end user and of all information on his/her account.

Network Management and Control Systems

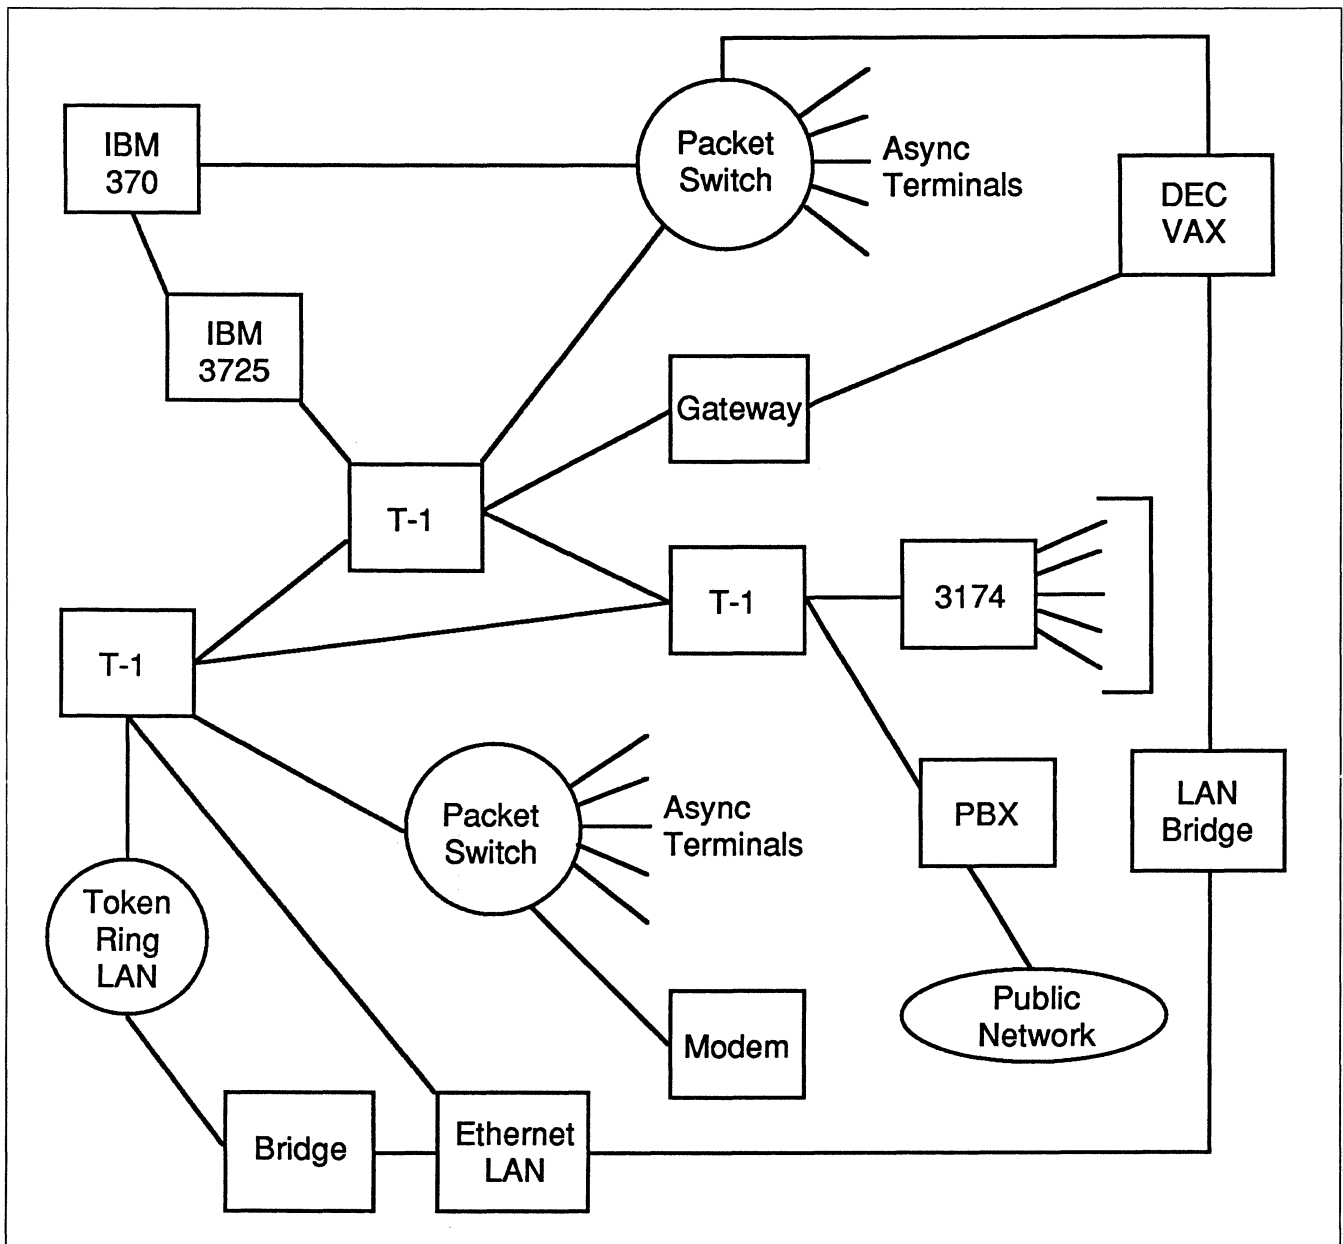


Figure 1. Typical private communications network.

Interface Requirements

The NMCS must be capable of interfacing with all of the systems that are part of the total communications network. For example, in Figure 1 is depicted a typical communication network that operates within an IBM SNA environment. The front end processor, 3725, interfaces with a private network that utilizes a T1 backbone. The backbone is part of a separately controlled network and the lines to the remote 3270 cluster controller are submultiplexed onto the overall T1 network. The specific interfaces are:

- The IBM SNA environment.
- The DEC DECnet Environment.
- The local RBOC terminations at each end. These may be in separate regional areas and thus may require separate control interfaces.
- The T1 network managers and muxes.
- The local modems.

Network Management and Control Systems

- The PBX switch through which the 9.6 Kbps lines are switched.

This is an example of how it is necessary to include all of the elements that are part of the total switched network as well as the dedicated network.

ARCHITECTURAL ALTERNATIVES

The architectural alternatives for the NMCS systems are based upon the need to satisfy the end user requirements. The environment for the NMCS is that of a computer communications network and far exceeds that of the standard voice or even data communications network. The computer communications needs are structured along the lines of the OSI layers and the NMCS must be designed to support those layers separately and in unison.

OSI Layers

The NMCS problem is one that encompasses all of the OSI layers. Most NMCS systems at the present are not layered in the OSI fashion but it will become quite clear that such layering elicits a clear understanding of all of the functions that have to be performed and where such functions are to be best performed.

The seven OSI layers and their NMCS functionality are best described as follows:

- Physical: This is the lowest level and is the one that the BOCs have a long history of addressing. The need at the physical level is to ensure that connectivity of the circuits is maintained.
- Data Link Level: At this level the issue is the ongoing point to point connectivity of the modems in the network. Typically, this is provided by the loop-back testing of the modes as well as the ability to utilize datascoopes for protocol testing. A typical problem in SNA networks may be the setting up of SDLC framing sequences and the ability to provide the correct bitframing. In addition, with multiprotocol networks, those running SDLC, BISYNCH, and X.25, the NMCS must have access to all of these remote protocols.
- Network: This is the point-to-point addressing issue in the network. Typically in a packet network, the control access to these points is in the network providers' PADs and is not readily accessible to the NMCS. In more sophisticated networks using T1 switches, control may be available through the switch control ports.

- Transport: This end to end addressing capability allows for the total integrity of the signal. It is typically the purview of the CPU and its associated front end processor. In an IBM SNA environment, the control is in NPDA and NCCF.
- Session: As with the transport layer, the control is at the CPU site and typically resides in the same locations for SNA as do the transport functions.
- Presentation: At the present time, the presentation layer is not supported as effectively as the other layers. With increasing complexity of presentation formats, there is a need to expand this capability. For example, the use of X.400 type formats will require a closer control of this layer. A single terminal, for instance, could have multiple formats used for presentation and these must be supported.
- Application: This layer is not supported at all in most NMCS systems. In particular, if the end users find themselves in a problem area, that problem may be in the application layer only, and not in any of the other layers. It will be essential for the NMCS to monitor and have access to this layer.

Functional Elements

The functional elements of the NMCS architecture include three major areas: software, hardware, and interface integration. These functional elements are the same for all of the seven layers in the design and change only in the specific implementation at the given layer. The software functions of the architecture include the following:

- Operating System: At this level the requirements are those of the platform that is being used as the controller for the appropriate layer. In some cases there may be more than one due to there being more than one controller. The level also requires that the files and Database be compatible.
- Support Services: This level of the software functionality provides the necessary generics to ultimately support the end user requirements. It may be viewed as a support shell to the overall system.
- Applications: These are the end user specific functions in the software that focus on layer specific support functions.

The software is typically layer specific and in most cases vendor specific. There is a communications interface that allows for the monitoring and control of the network subelement. It is generally through this

Network Management and Control Systems

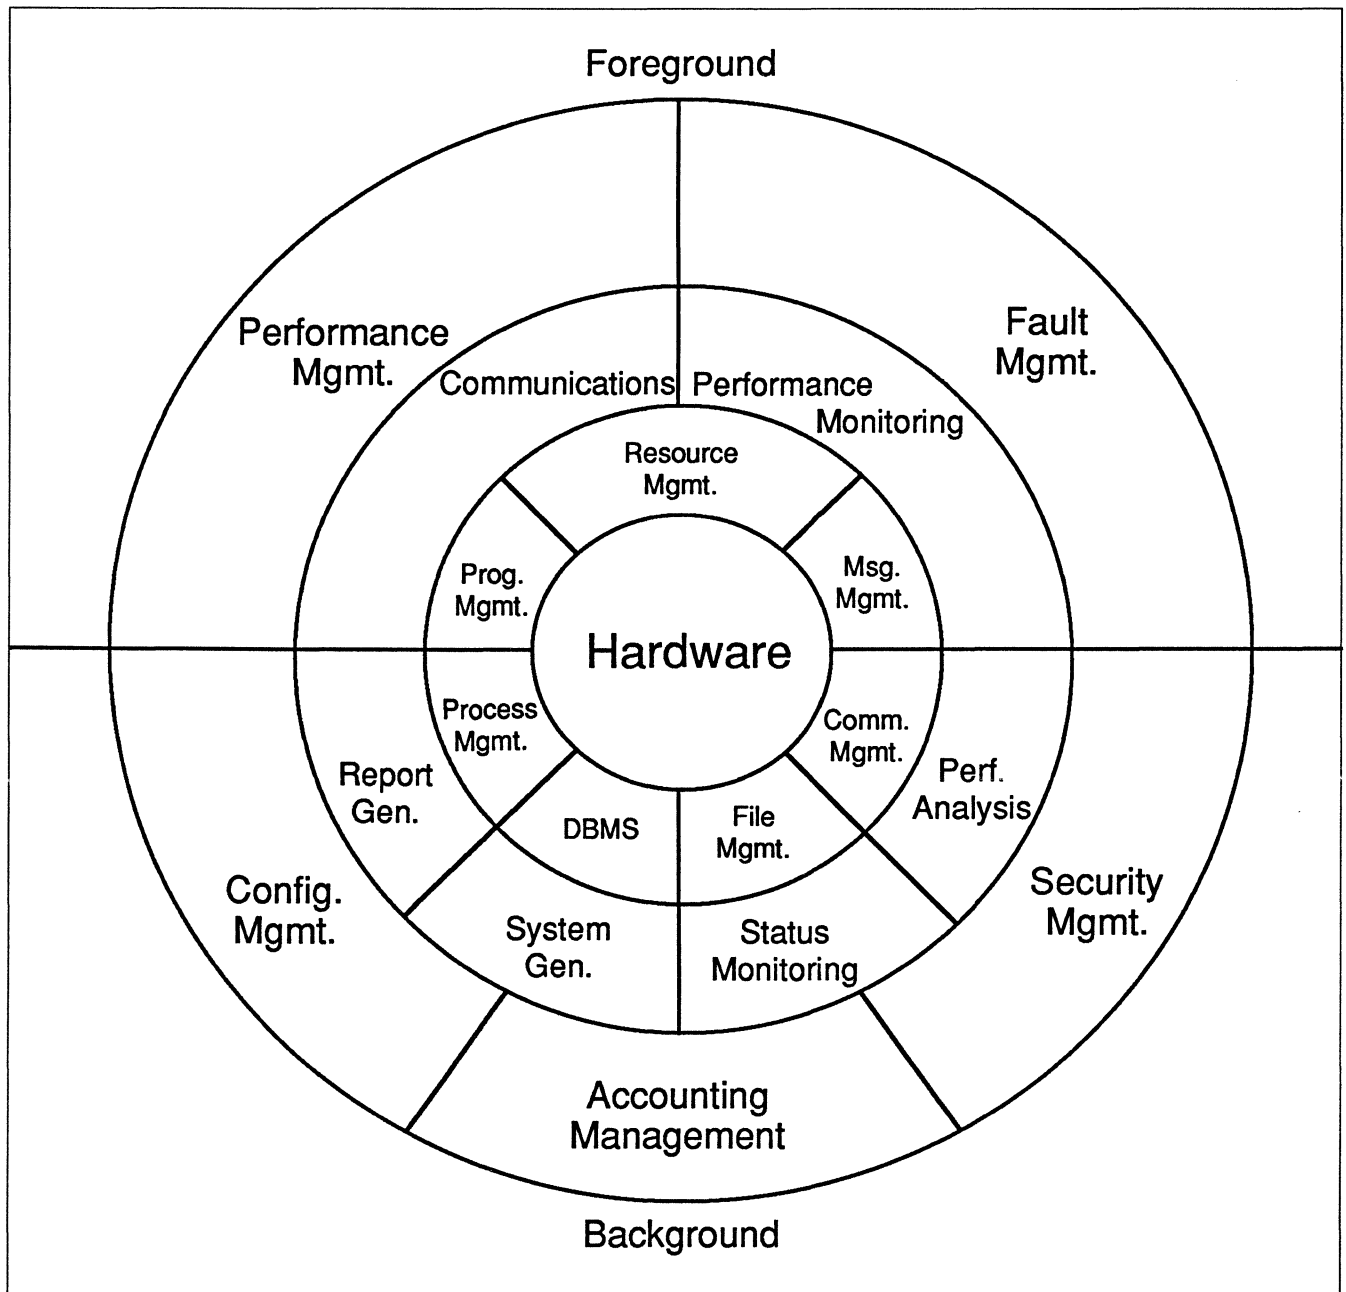


Figure 2. Network management software architecture.

port that the manager of the network will have access to the local software elements.

The hardware functional elements are not as structured as the software. In many cases the vendor specific equipment has a generic set of NMCS functions resident in its own hardware. In other cases the hardware for NMCS is a separate machine.

Figure 2 depicts a software architecture for a general NMCS system. In this architecture we have shown all of the three major software layers. The functions are

generic and they must span not only each physical element in the network but also each layer in the OSI hierarchy. This latter requirement is a critical factor in delivering effective network management.

Layer Interfaces

In order to support the end user requirements, the seven layers of management must be interconnected in some fashion. In the present environment, the interconnection in many cases is by human interac-

Network Management and Control Systems

tion. There is a clear need to automate this function and to allow the end user to have access to that single point of problem resolution. Thus the layer interfaces must eventually adhere to a single set of standards. At present however the interfaces are nonstandard. In particular, the interface at any one layer is different and may be nothing more than an RS-232 interface to a monitor. Thus it is necessary to provide for support of this minimal interface.

Multidomain Operations

In an NMCS environment there are multiple domains of operation. We define a domain as a partially enclosed environment that supports one host processor and concentrates on one level of the OSI domain. For example, in an IBM and DEC combined environment there may be seven domains for each host, for a total of fourteen domains. In this environment it will be necessary to provide the functional interface between the domains and to allow for the interexchange of information between the domains.

Consider the network configuration depicted in Figure 1. Here the network includes an IBM machine, T1 network controllers, LANs, Bridges, and a DIGITAL VAX. The intention is to provide the user with the ability to access the control ports of all of the terminals and to allow for a single and unified control station for the system. In this case the user has two computer domains and is interested in the control of 6 layers in each domain.

Domain Interfaces

The domain interfaces can range from the simple to the complex. This is best described as follows:

- **Independent:** The user has a single screen for each of the elements in the network. This screen allows for the remote accessing of the managers for each of the separate network elements.
- **Interconnected:** In this architectural alternative the terminals are not just duplicated physically, but a single software shell interconnects the different control screens into a single effective system. The separate control formats are maintained however and the user is still using the separate control ports as was done in the independent case. However, the single shell interconnects the screens so as to allow the user to do this from a single terminal.
- **Integrated:** In this case the shell is more encompassing and is supported by a set of intelligent drivers that allows the controller to exercise commands to

each of the network elements in a consistent format. The controller does not need to learn a large set of different diagnostic and control commands but only a single set. This allows for the rapid redeployment of network resources and dramatically reduces the cost of network management.

To effect these types of architectures, the resources vary drastically. The first approach is merely a remote monitoring and control of devices that are naturally part of the existing elements. The second approach allows for the signals to be transmitted and received as if they were separate monitors, but they are transformed at the presentation layer into a single screen. The last approach requires the development of separate drivers to support the different variety of network elements.

SYSTEM INTERFACES

SNA

Any NMCS system must support the interface to the IBM SNA devices, both physical and logical. With the introduction of NetView, IBM has enhanced this access and thus will allow the user to utilize more sophisticated interfaces that will integrate many devices and ultimately, using the intelligent driver approach, will allow for a totally integrated system. In particular, the IBM features of NPDA and NCCF will be essential in terms of supporting an IBM network.

ISDN

In the ISDN configuration, the end user will have access to the D signaling channel as well as other potential network services. The signaling channel, enhanced with the capabilities of Signalling System 7 will allow for the direct connection with the network control functions and will provide the basis for the interconnection of the Interexchange Carriers (IC) to provide data on the Data Link and Network layers.

Multivendor Interfaces

The development of the recent OSI standards on network management provide guidelines to other vendor manufacturers that allows them to focus their attention on the development of consistent interfaces and data formats that will allow the support of fully integrated systems.

Network Management and Control Systems

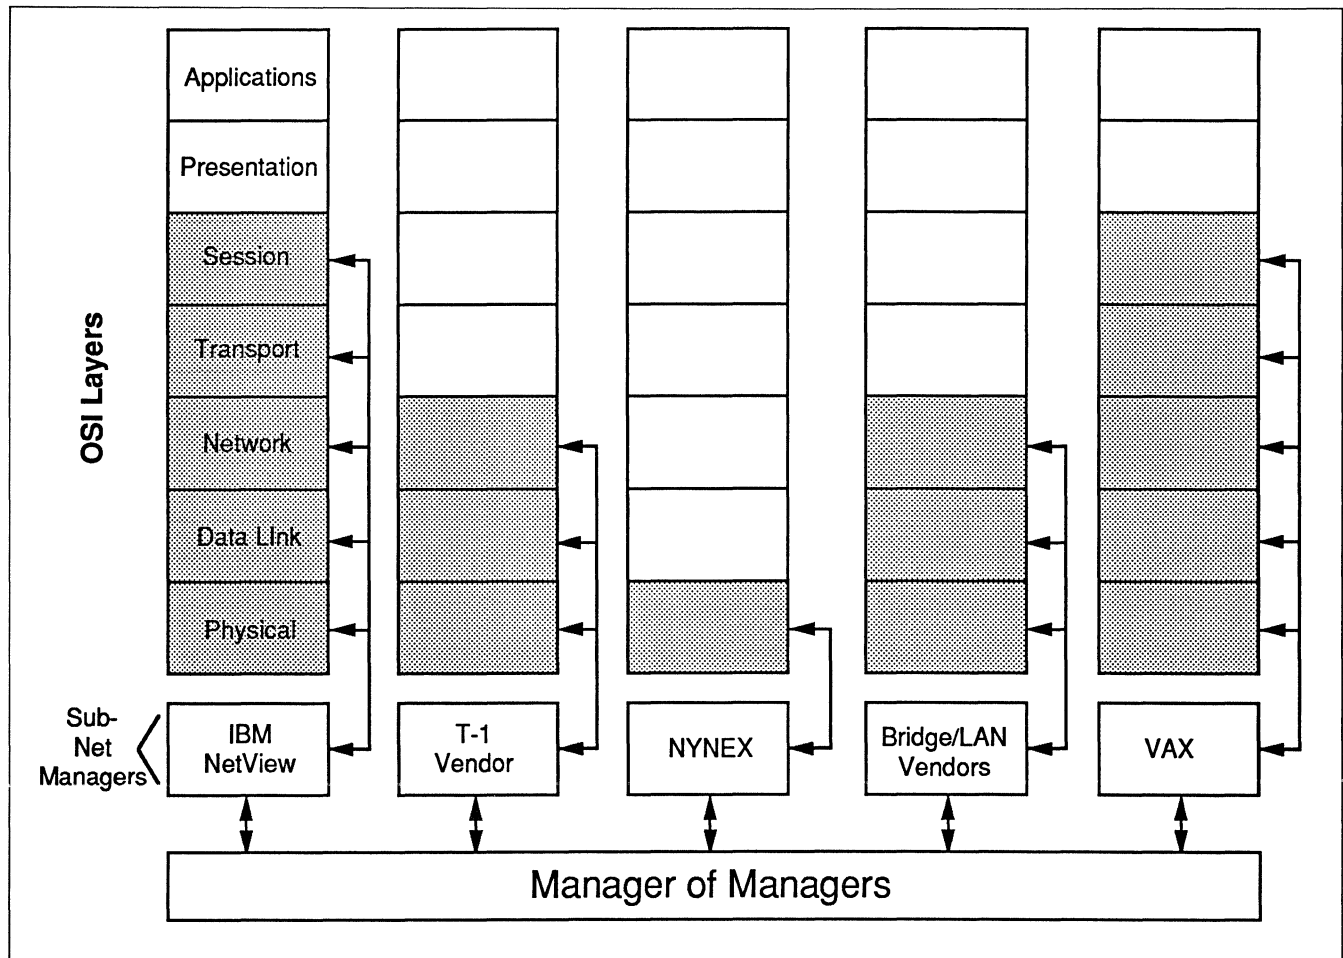


Figure 3. Network management integration concept.

NMCS FUNCTIONS AND ELEMENTS

The NMCS has functions that are performed in both a foreground or near real time mode and a background or non-real time mode. The generic sets of NMCS functions are common amongst all of the OSI layers that have been described. These functions may be complete at a layer or may be limited in scope. In the RBOCs, the physical layer includes all of the functionality described. The interlinking of these functions, even within the physical layer, presents a significant task. As a figure of merit, the total lines of code in the NMCS for the physical layer exceed 50 million lines.

The layers structure of the NMCS is described in Figure 3. For each level the functionality is generically the same. The ability to interconnect the layers is done through the communications layer.

The ideal NMCS architecture is one which allows the NMCS to act as the manager of managers. This concept provides for the support of all the functionality

through interfaces to the individual subnet controllers and allows for a common presentation and control format to the network manager.

Foreground Functions

The foreground functions provide direct support to the near real time operations of the NMCS. Specifically:

Network Resource Management

This function provides for the real time network reallocation of resources. For example, we can describe how this function may be applied at each layer:

- Physical: This will allow the reallocation of twisted pairs and fiber backup. It permits the alternate routing over different physical media.

Network Management and Control Systems

- **Data Link:** This provides the alternate switching between modems in the network.
- **Network:** This permits the alternate routing in a packet network. In this case the NMCS function may be part of the protocol used in the PADs.
- **Transport:** At the CPU during an end to end outage, the CPU through NCCF may be able to reroute to an alternate host in a multidomain SNA session.
- **Presentation:** When a user logs on with a different application, switching from 3270 emulation to VT 200 emulation, the NMCS may be able to in real time, download a new presentation format controller.
- **Applications:** At this level, there is the need to allocate different applications programs and this function is typically performed by the data base administrator.

Problem Determination

This foreground function provides for the real time determination of problems at each of the network layers. It collects data from the performance monitor and, using sophisticated algorithms, provides a determination of the problem, its cause, and a possible set of corrective measures.

Communications

The need to communicate is twofold. First the system must be able to communicate in its own layer. This communication is between the NMCS at that layer and the elements that are being effected. Second, the NMCS must be able to communicate between the layers and the NMCS functions that transcend the entire network process. The communications function is robust to the overall needs of the network.

Performance Monitoring

The performance monitoring function is a data gathering function that acts as an input to the various elements that work as a whole in that layer. For example, the performance monitor at the Data Link level gathers data from the modems on the error rate, delays, and queue sizes. This gives a performance measure on the protocol transmission and throughput.

Background Functions

The background functions are the most complicated functions and are those that are most often done in a manual form and may be left as the last to be implemented. As has been observed in many networks, the background functions are most often associated with the back office operations. The integration of these functions will assist in the full utilization of the network.

Network Configuration

The function provides the NMCS with a full inventory of what is configured as what in each level. For example the Transport configuration is set in a table in VTAM in an IBM environment and is accessible through NetView.

Traffic Management

The function provides for the management of the traffic and the connectivity of that traffic in the network. At the lowest level it is comprised of the physical links in the network and at the highest level it is comprised of the utilization of certain applications programs and data bases.

Administration

This includes all of the standard back office functions that are performed at the appropriate level.

Customer Service Interface (CSI)

The most important element in an NMCS is the ability to assist the end user with network problems. The CSI provides to the Customer Service Representative (CSR) a means to identify the end user, to determine the problem, to provide for problem resolution and to ensure that restoral of layer service has occurred. In most existing networks, the end user has no access to a CSR and typically has to negotiate through a maze of intermediaries. In the fully interconnected NMCS, the CSR will assist in all functions.

Performance Analysis

The performance analysis function provides for the interfacing of the foreground data gathering functions and developing and displaying the results of the net-

Network Management and Control Systems

work performance. This allows for an ongoing level of quality of service and the visible monitoring of that service.

Billing

A billing function may be required at each layer. This is too often forgotten except as part of the common carrier networks. As part of any network and integral to the NMCS the utilization and direct billing or expense allocation portion is a necessary element.

EVOLUTION

Status Now

The present state of NMCS systems is highly fragmented. In the IBM world the user is seeing a proliferation of NetView and its options. Each separate vendor offers a network manager of some type and none of the vendors provide compatibility. The access to the networks that provide transport is sketchy if at all existent. The controllers are all independent and require significant training.

The managers of networks typically have a collection of network control devices and are continuously trying to keep up with the most recent releases of the software.

Desired State

The desired state of evolution is to have a manager of managers (MOM) that keeps the network and all of its elements working effectively. The network manager must recognize the needs of the end user and integrate into all of the OSI layers.

The manager of managers approach consists of the following elements:

- **Interface:** A common set of interface protocols that allows the MOM to interconnect with all subnet managers in a standard format.
- **Control:** A common command language that allows the manager to control all subnet elements from a single terminal control point. The common language must be capable of both query interaction and menu interaction.
- **Presentation:** The MOM must provide the manager with a single integrated set of information from each of the subnet controllers and should allow the manager to reformat the data elements in a single fashion to modify the presentation for a specific

application. This is typically driven by the need of each network to have custom driven management tools.

- **Interconnect:** The MOM must have integral to its operation the ability to interconnect with the customers other back-office system and support information to provide database support and report generation. All too often the approach is to provide a tool that satisfies the needs of the moment but neglects the customer's needs for ongoing support and growth.
- **Problem Determination:** The MOM concept should have the internal intelligence to anticipate, identify, isolate and circumvent network problems. This means that there must be some form of Artificial Intelligence that adapts to the network's performance and assists the network manager in performing his or her tasks.

This desired state for NMCS architectures is evolving and there appears to be no one approach that meets the needs.

Evolution Path

The path to get from where we are to the ideal state should include the ability to integrate and to interoperate with diverse elements. Thus, it will be critical that the providers of network management elements for subnetwork elements recognize the need for the interfacing of their products in a global fashion and to assist their customers in developing an integrated NMCS system.

CONCLUSIONS

This report provides an overview of the NMCS problem and a view of an architecture that is evolving in the future. The NMCS architecture for the future, especially in the world of computer communications networks, will require the ability to communicate and control all of the seven layers of the OSI model. It will require that the front end and back end systems be integrated and that the end user be the focal point of operations.

There is an evolution in the area of NMCS, and that evolution is to a totally integrated package of services. The goal is to provide the end user with a transparent connection to a support infrastructure that allows seamless service. However, with the proliferation of new network elements, the needs are directed at NMCS architectures that allow for the integration of dissimilar elements. This report suggests such an architecture. □

IBM's Approach to Network Management

This report will help you to:

- Evaluate IBM's strategic product for network management—NetView.
 - Use NetView components for operations, problem, and performance management.
 - Examine NetView's ability to manage both SNA and non-SNA systems and products.
-

IBM's approach to network management involves two related concepts. The first of these is to provide a family of products that allows enterprises to manage their networks from a centralized point (Figure 1). This family offers the major disciplines of network management—operations, problem management, change management, configuration management, and performance and accounting management. The NetView program is an important member of this family.

The second concept is to use distributed points of control throughout the network to support both Systems Network Architecture (SNA) and non-SNA systems and products. These points of control are used to collect network management information, forward it to a central site, and act as operations receiving points for commands coming from the central site.

These distributed points are provided by the open network management functions of NetView, such as the Service Point Command Interface (SPCI) working in conjunction with NetView/PC¹ or user applications.

The centralized and distributed points of control are tied together by IBM's SNA-based network management architectures. As an example of our open management direction, the NetView program provides the capability for all products, IBM or non-IBM, to contribute to the problem management function by providing a generic architecture to embody the alert structure.

At the foundation of the network management concept is the NetView program.² NetView is an integrated product formed by consolidating five previously available IBM network products. There are four major components of NetView: Command Facility, Session Monitor (including support for the IBM 4700 Communication Finance System), Hardware Monitor, and Status Monitor (Figure 2). These components are combined with usability and installation enhancements such as help, browse, and installation aids. The components work together to provide an enterprise with a common operations interface to

This Datapro report is based on "An Integrated Network Management Product," by Denise Kanyuh, pp. 45-59. © 1988 by International Business Machines Corporation. Reprinted, with permission, from the *IBM Systems Journal*, Vol. 27, No. 1.

Index to This Report	Page
NetView Components	103
Problem Management	108
Open Network Management Facilities	110

IBM's Approach to Network Management

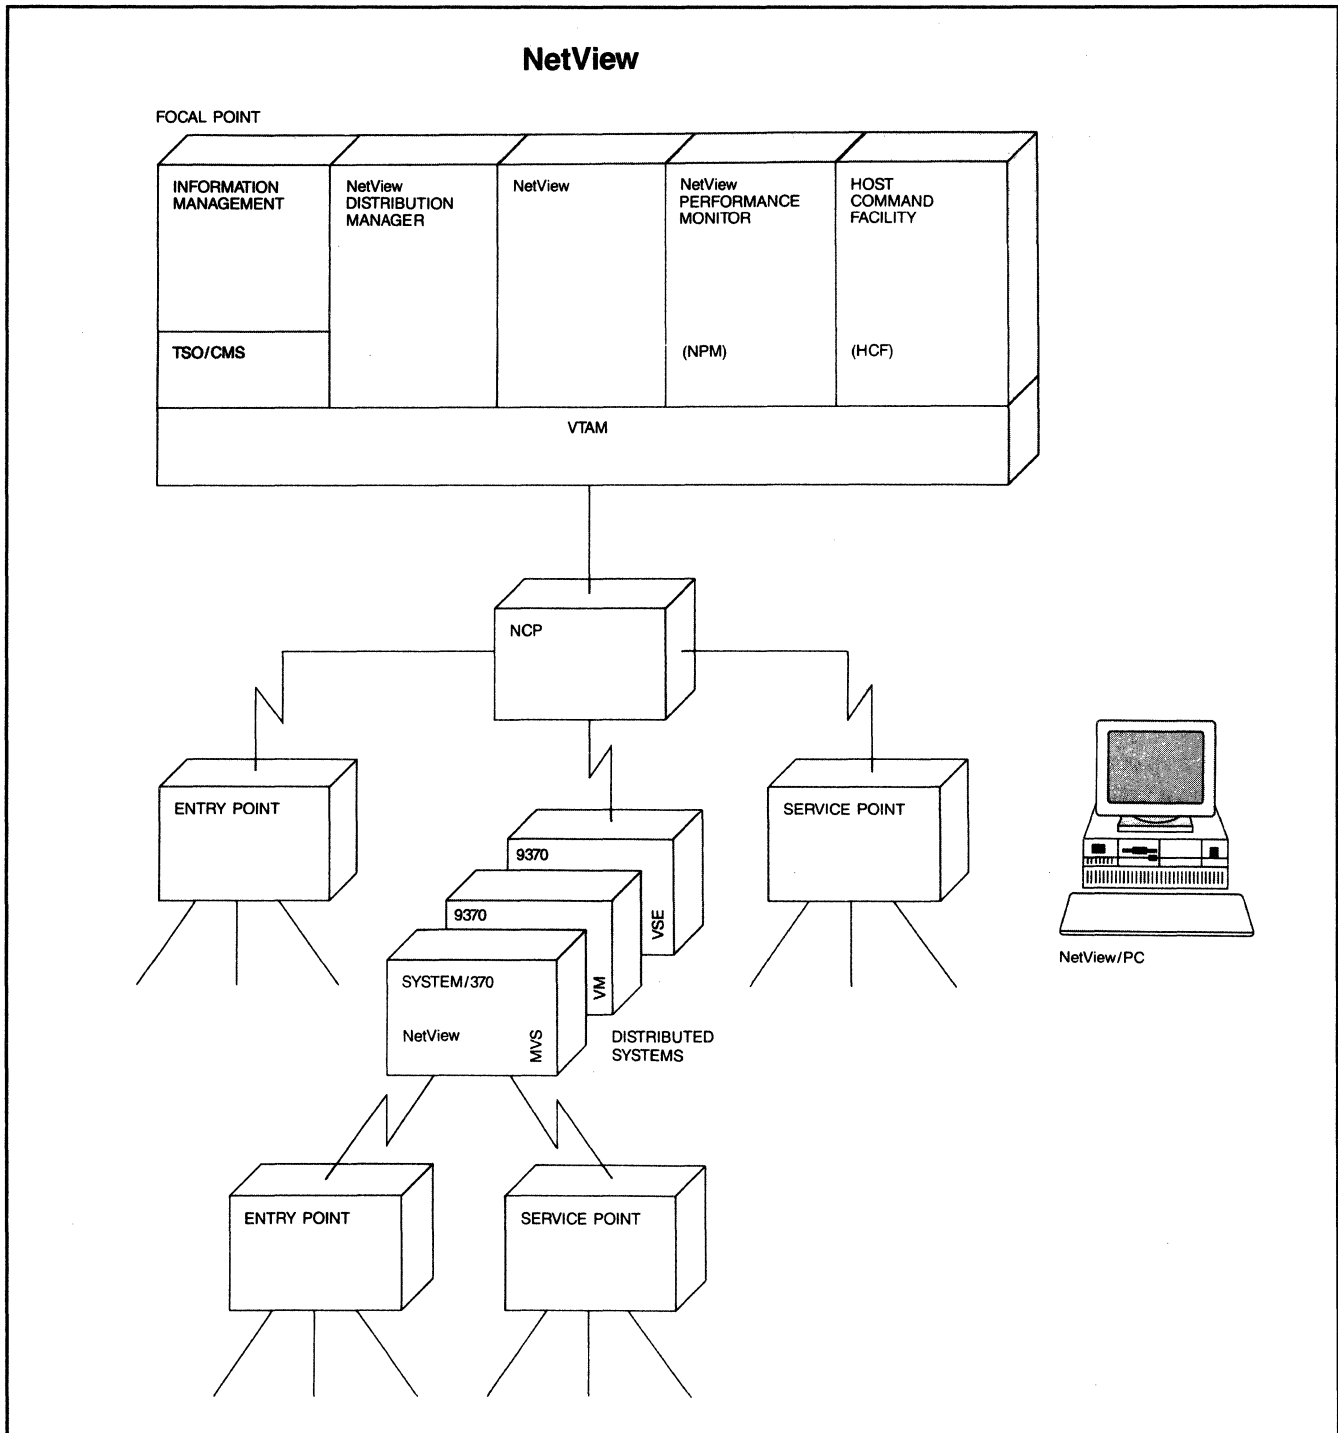


Figure 1. IBM's network management system products.

manage and automate its network, to aid in problem determination, and to improve performance. NetView has extensive facilities for open network management both at the focal point and in communication with the distributed points of control.

The motivation and method behind the formulation of NetView will be discussed next.

IBM's Approach to Network Management

THE ORIGINS OF NETVIEW

The largest and most important force behind the formation of the NetView program was the user—in particular, IBM's networking customers. Not only was the initial decision to develop NetView a direct result of what the marketplace wanted, but major design points were also based on users' input and feedback.

The need to simplify the installation and operation of our network management products came to us from three independent groups of users: a communication network management marketing task force, the Communications Programming Customer Council, and the user organization, SHARE. In a SHARE White Paper,³ a detailed account of the product integration was given, including the need for help facilities and a common presentation services/user interface.

The decision to combine the products was only the beginning of the design. Before the NetView program was developed, a usability test was run on the network management programming products Network Logical Data Manager (NLDM), Network Communication and Control Facility (NCCF), and Network Problem Determination Aid (NPDA), along with two field-developed programs, the Network Management

Productivity Facility (NMPF) and VTAM Node Control Application (VNCA). This test, in which users participated, identified several problems with the existing product set. Some of the major areas of concern were consistency of user interfaces from product to product, including program function (PF) keys, colors, and command input areas, and ability to move easily from one product to another. NMPF include a set of tutorials and help desk scenarios which the participants overwhelmingly agreed were an asset.⁴

With this user input, the development of NetView commenced. Once a direction was set, the concepts were reviewed with over 60 customer accounts. This was done through customer calls and surveys, FOCUS sessions, and the customer council. These sources agreed with the direction that the merged product was to take.

The first release tackled the external interface problems and the interactions between the products, now components. The installation process was greatly simplified, reducing the installation time and likelihood of errors. An extensive help facility was incorporated into the product as well as a help desk offering. Release 2 used this work as a base to answer another pressing requirement: the ability to automate network operations.

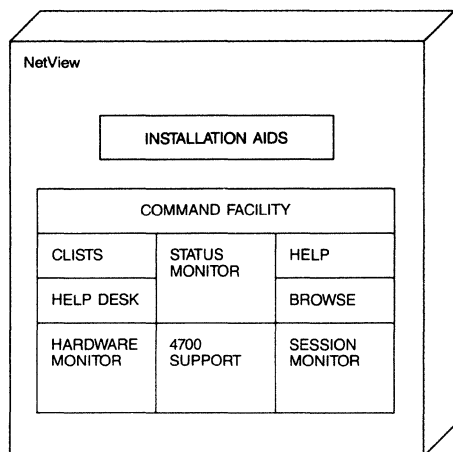
A second usability test with users was run, this time on NetView Release 1 before it was shipped. When the results of this test were compared with those of the previous test, it was shown that the NetView program was significantly more usable than the set of products preceding NetView.⁴

The following sections will explain the structure of NetView and the areas of network management that it addresses.

NETVIEW COMPONENTS

Command Facility. The Command Facility is the base component of NetView.^{5,6} Through this component services are provided for centralized network management. These services include message routing, logging, presentation, and automation. Operator support is also available from the command facility. Commands or CLISTs (command lists—lists of NetView commands incorporated in a simple language facility) entered by the operator can be routed to different domains or systems within a domain.

The other NetView components also rely on the Command Facility as their program base. Macro services, including database access, operating system independence, intracomponent message routing, and



COMMAND FACILITY
—BASE FOR CENTRALIZED NETWORK MANAGEMENT
—NETWORK OPERATOR SERVICES

SESSION MONITOR
—RESPONSE TIME MONITORING
—SESSION CONFIGURATION
—DIAGNOSTICS

HARDWARE MONITOR
—PROBLEM ALERTING
—RECOMMENDED ACTIONS
—PROBLEM DETERMINATION DATA COLLECTION
AND DISPLAY

STATUS MONITOR
—NETWORK STATUS DISPLAY
—AUTOMATIC REACTIVATION

Figure 2. The components of the NetView program.

IBM's Approach to Network Management

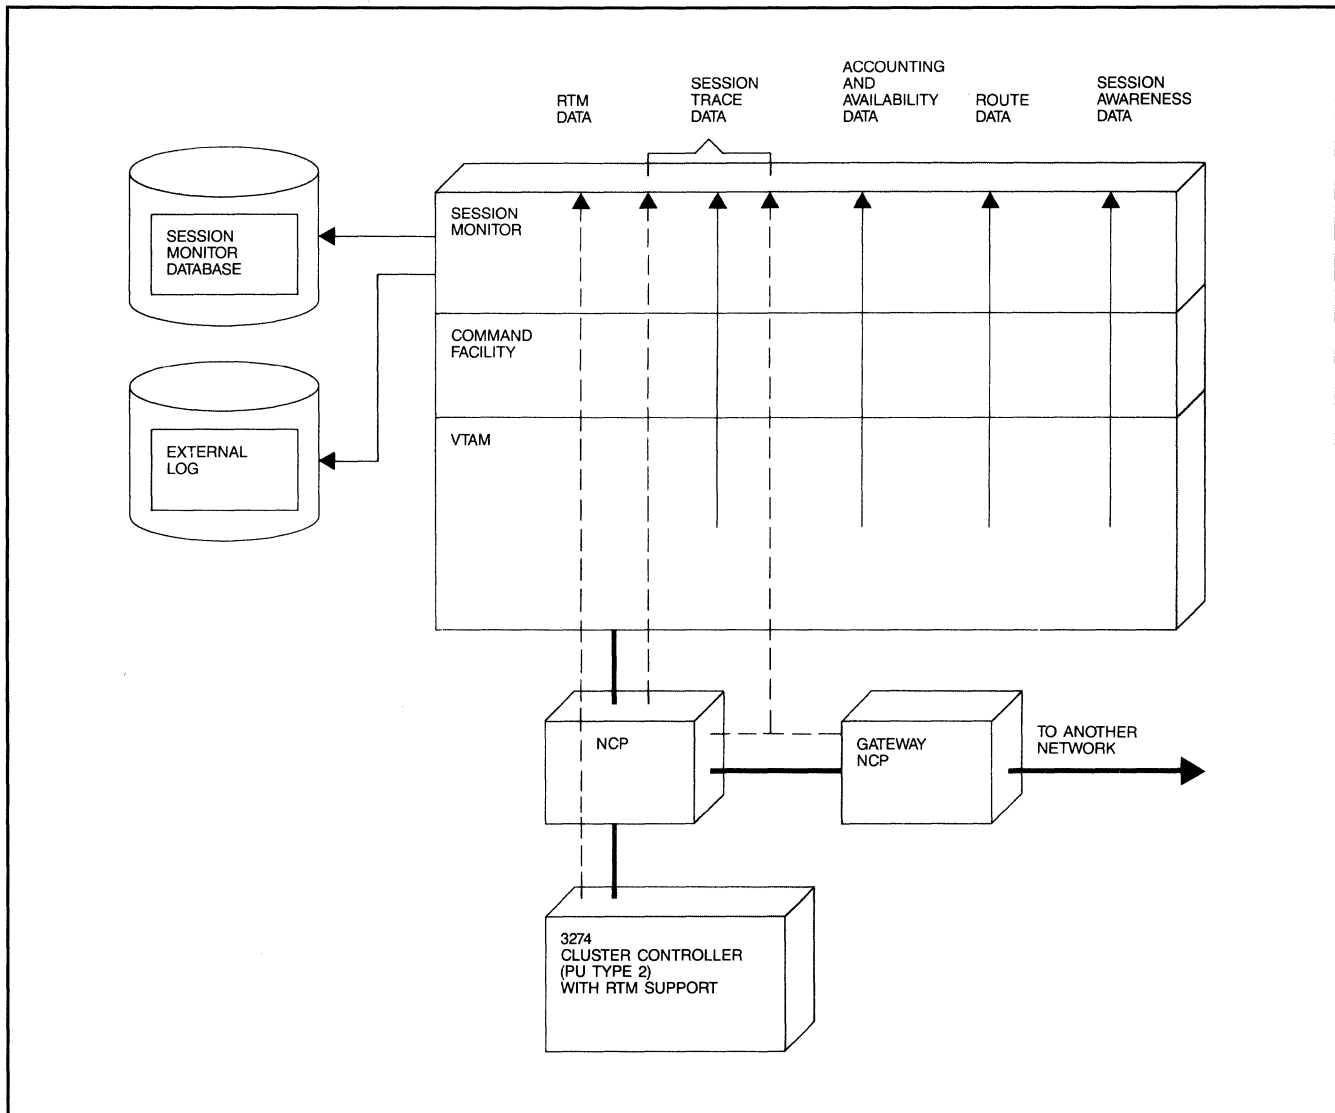


Figure 3. Session monitor data collection.

screen handling and presentation, are available. The Application Programming Interface (API) for user-written customization also interfaces with the Command Facility.

The Command Facility is the main contributor to the operations discipline in NetView. This function will be explained in more detail later in this report.

Session Monitor. Diagnostic facilities for the SNA logical network sessions are provided by the Session Monitor. This component collects and correlates data about SNA sessions and provides on-line, interactive access to the data. Through this function, logical network problems and error conditions can be identified in a productive manner. These conditions include “hung” sessions, lost messages, and route problems.

The Session Monitor collects data about same-domain, cross-domain, and cross-network SNA sessions. The following types of data are collected for these sessions (Figure 3).

- Session response time data—This information is measured and accumulated by control units with the Response Time Monitor (RTM) feature. It is sent to the Session Monitor on request.
- Session trace data—Such data consist of session activation parameters, Virtual Telecommunications Access Method (VTAM) Path Information Unit (PIU) data, and Network Control Program (NCP) data.
- Network account and availability measurement data—The Session Monitor provides data on net-

IBM's Approach to Network Management

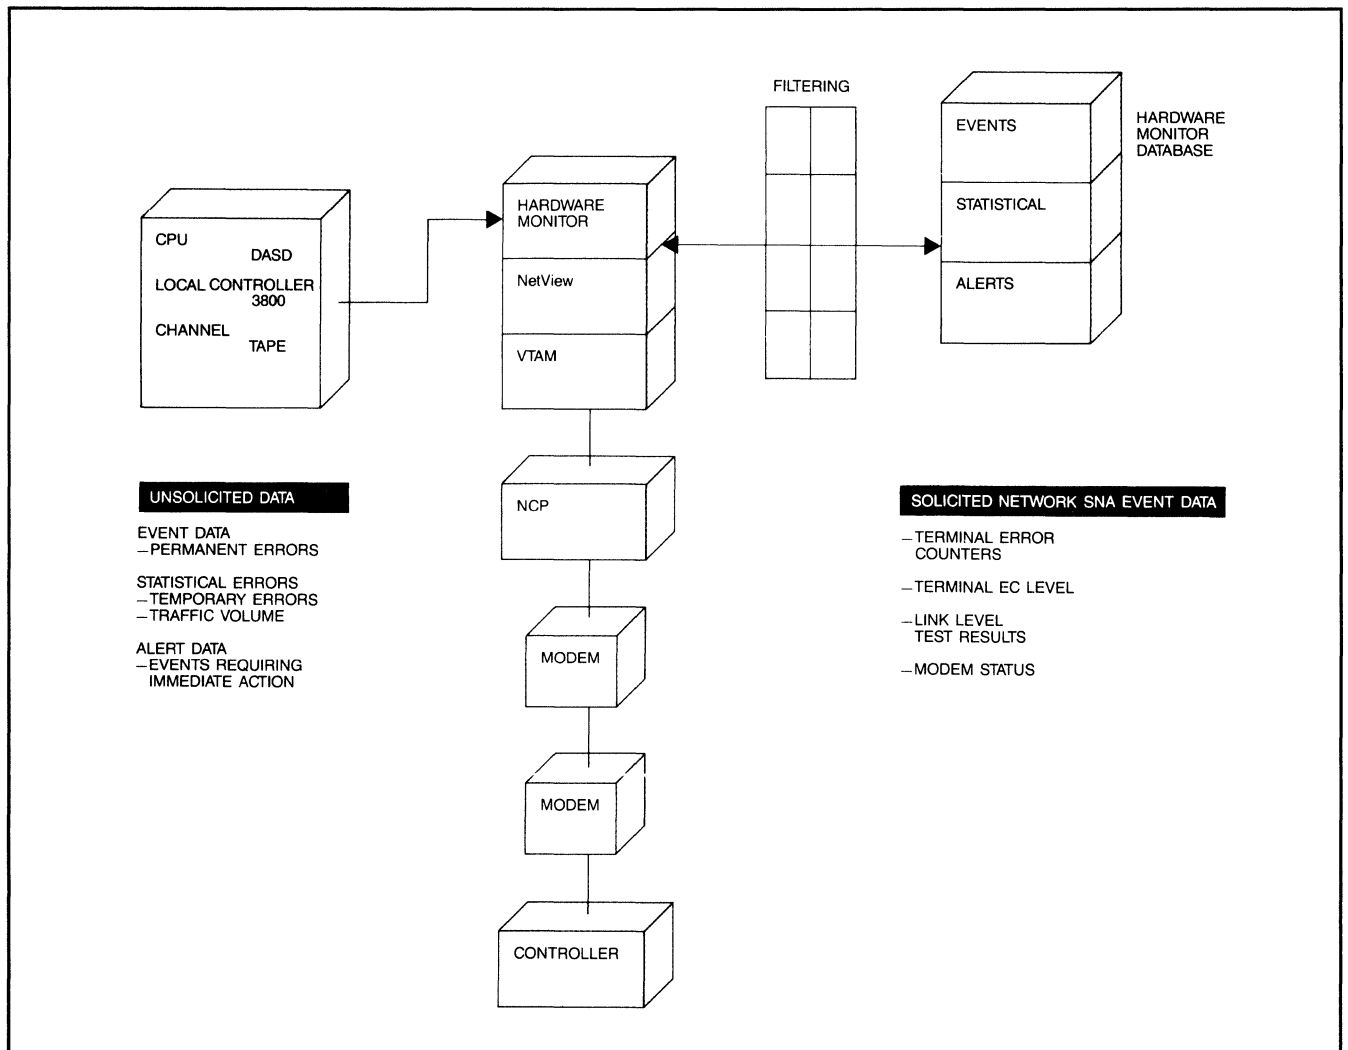


Figure 4. Data collected by hardware monitor.

work availability and distribution of usage of network resources. These data are written to an external log.

- Route data—The list of Physical Units (PUs) and Transmission Groups (TGs) that make up an explicit route are included here.
- Session awareness data—The data include information about session activity, identifying the session partners, and configuration.

Hardware Monitor. The Hardware Monitor collects error and information records from network devices (both SNA and non-SNA). It provides interactive presentation and logging capabilities for these records based on customizable filters. These functions assist network personnel in performing network problem determination. These data are also analyzed for probable cause and recommended actions, which can iso-

late a network problem to that of a specific failing component. The alert feature informs the operator quickly of high-priority problems.

The data collected by the Hardware Monitor can be classified as solicited or unsolicited (Figure 4). Solicited data are received as a result of a specific request for information. Unsolicited data, such as the result of an error being detected or counters exceeded, are received without any action.

The records that are sent to the Hardware Monitor can be one of three types:

1. Statistics—These are records of traffic and recoverable error counts.
2. Events—Unexpected occurrences in network operation cause event records such as resource activation failure to be generated.

IBM's Approach to Network Management

STATMON.DSS		DOMAIN STATUS SUMMARY						05:48
HOST: HOST11		*0*	*1*	*2*	*3*	*4*		
		ACTIVE	PENDING	INACT	MONIT	NEVACT	OTHER	
....4	NCP/CA MAJOR4	
..121	LINES	...8732	...2	
..226	PUS/CLUSTERS	...16	...32178	
..937	LUS/TERMS	...88	...22	...3824	
...9	SWITCHED MAJ	...9	
...12	SWITCHED PUS12	
...44	SWITCHED LUS9	...35	
...2	LOCAL MAJ NDS	...2	
...30	LUS/TERMS	...25	...14	
...19	APPL MAJ NDS	...19	
..378	APPLICATIONS	...3039	...309	
...1	CDRM MAJ NDS	...1	
...17	CDRMS	...7	...73	
...25	CDRSC MAJ NDS	...25	
..100	CDRSCS	...981	...1	
-----		-----	-----	-----	-----	-----	-----	
..1925	TOTAL NODES	..411	...62	...31090	..359	
CMD-->								
1-HELP 2-END 3-RETURN 4-BROWSE LOG 6-ROLL				9-REFRESH				

Figure 5. Status monitor summary panel.

3. Alerts—Events that require immediate attention are alerts. Alerts are determined by the filters defined.

Status Monitor. The display of the Status Monitor allows the user to tell “at a glance” the status of the components in his or her domain (Figure 5). This status information can then be used to control these resources.

The Status Monitor uses colors to identify the different states in which a resource can be. For example, active nodes are presented in green, inactive in red, and pending in pink.

From the main status summary screen a series of detail panels can be displayed. These panels allow the user to view a subset of the resources (such as all inactive Logical Units, or LUs) and include more information on the selected resources. This information can be node names and description, summary of node status over time, or message traffic counts to and from an application program node.

The Status Monitor provides another function, the monitoring of inactive nodes. This component has the capability to reactivate minor nodes that are in the inactive state. These nodes appear in the MONIT column on the display. Once a node is reactivated, its lower nodes, if present, will be monitored and reactivated as well.

Usability Features. Working in and around the component structure of the NetView program are several features which enhance its usability and cohesiveness. Included in this area are browse support, extensive help facilities, a unique “ROLL” approach to command entry, and packaged installation aids.

A full-screen browse capability is provided for on-line viewing of NetView installation files, CLISTs, panels, and the message log. The browse panels provide a scroll field, program function keys, and a find command.

The message-log browse has an additional “Important Message Indicator” capability. Messages are user-defined to fall into one of four categories that can set off indicators on certain NetView panels. The operator is then led directly to the message in the log that set off the indicator. These messages can also be color-coded for easy visibility.

Operator assistance is provided by the on-line help facility when using NetView. Helpful information can be viewed for every NetView command, component, and certain displayed panel fields and codes. A step-by-step approach is used in the help desk to simplify network problem determination. Complex procedures are reduced to simple steps that lead the operator toward resolution of a specific problem. Diagnosis assistance covers NetView, VTAM, and other network offerings such as the IBM 4700 Finance Communication System.

IBM's Approach to Network Management

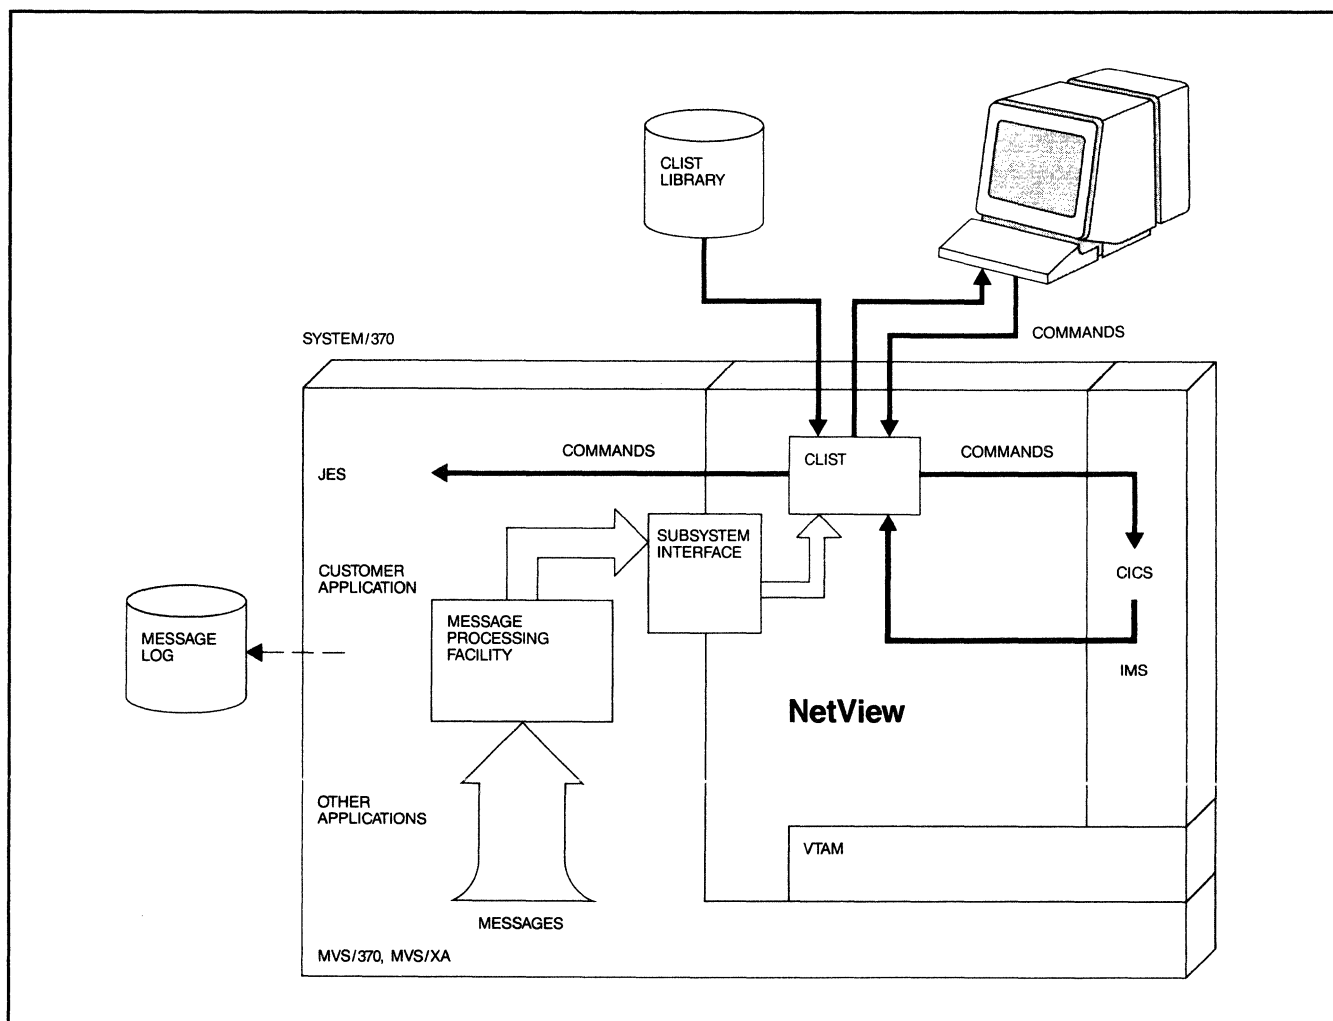


Figure 6. NetView system interfaces.

Through the use of a command line on all NetView panels and appropriate program function keys, a NetView command or component can be invoked from the screen of another component. Once a component has been entered, it is placed on a "stack." Using a program function key, the user can "roll" through the components on the stack, entering commands and changing screens as he or she goes.

NetView features an advanced installation process based on a prepackaged sample network and a task-oriented procedure. Together they simplify the user interface and reduce errors. The sample network, which is shipped in machine-readable form, includes VTAM, NCP, and NetView definition statements, along with the tables, CLISTs, and Job Control Language (JCL) to initialize the product. The installer then verifies that NetView is functioning properly by executing a predefined set of scenarios.⁷

Network Management Functions Provided by the NetView Program

Operational Control Functions. The purpose of the operational control functions is to provide the facilities and processes for controlling and managing all of the resources in the operational system environment. NetView, through the Command Facility, offers this capability to the network operator for managing both local and remote systems from a consolidated point or from distributed points of control. The operational control strategy is to minimize a human operator's activities through the use of automated operation. Extensive facilities in NetView allow this automation.

The first set of functions deals with the NetView operator. Security and control are provided by logon authorization, scope of commands, and span of control. The NetView operator must enter the proper identification and password to gain access to the system. A profile associated with each operator defines

IBM's Approach to Network Management




Figure 7. Extended remote operations and control.

what access to the commands, or scope, an operator possesses, and what span of control the operator exercises over resources. The scope capabilities limit an operator to a subset of network functions, and the span of control defines which resources an operator may control in the network.

NetView allows commands entered at its terminals to be routed to different system components, domains, or networks. The cross-domain and cross-network communication of NetView allows an entire network to be controlled from one operator station. As an example, with the Terminal Access Facility (TAF), an operator can log on to several subsystems, including the Time Sharing Option (TSO) and the Information

IBM's Approach to Network Management




Figure 8. Response time by collection period display.

Management System (IMS), without logging off NetView. TAF can also be used in cross-domain communication by logging on to a remote system that has NetView.

NetView is a focal point not only for networking commands but for system commands as well. Using the Multiple Virtual Storage (MVS) Subsystem Interface, NetView can process and route MVS system and subsystem commands issued from a CLIST or NetView operator (Figure 6). These commands can be routed to a subsystem in the same domain or to another domain. To make the function complete, an MVS operator can also issue NetView line-by-line commands and CLISTs from an MVS console.

Remote operations, including bring-up and restart capability for a target host, are possible through NetView and an additional offering, the Inter-System Control Facility (ISCF). The ISCF code in the host works with the ISCF code in an IBM Personal Computer (IBM PC) to allow commands to be issued from a central NetView control host and routed to a target host. These commands can be used to perform remotely those functions that would otherwise require manual intervention by an operator at every CPU location (Figure 7).

The automation facilities in NetView comprise mainly two functions: an automation task and a message automation facility. One or more automation

tasks may be active in NetView, combining with the message automation support to form a comprehensive set of capabilities.

All operator task functions are supported in the automation task except for full-screen presentation services, since this task has no operator console. This task removes the dependency of NetView on VTAM, allowing those functions which do not require VTAM, such as certain Hardware Monitor operations and message automation, to continue running throughout VTAM outages. The automation task has the ability to restart VTAM and to start the other NetView functions after VTAM returns from an outage. This support allows NetView to be brought up before VTAM is initially active. The most significant result of this VTAM dependency removal is that NetView may be used for system automation in systems that do not use VTAM.

The message automation and CLIST facilities of NetView can be used to control operator functions in multiple console support (MCS), Job Entry Subsystem 2 (JES2), JES3, NetView, and any application that can be accessed by TAF. A NetView CLIST can issue commands to and can intercept messages for processing from any of these systems.

The message table support allows for a variable number of criteria to be specified and compared against an incoming message. These criteria may range from string comparison in the text to the JES job identifier

IBM's Approach to Network Management

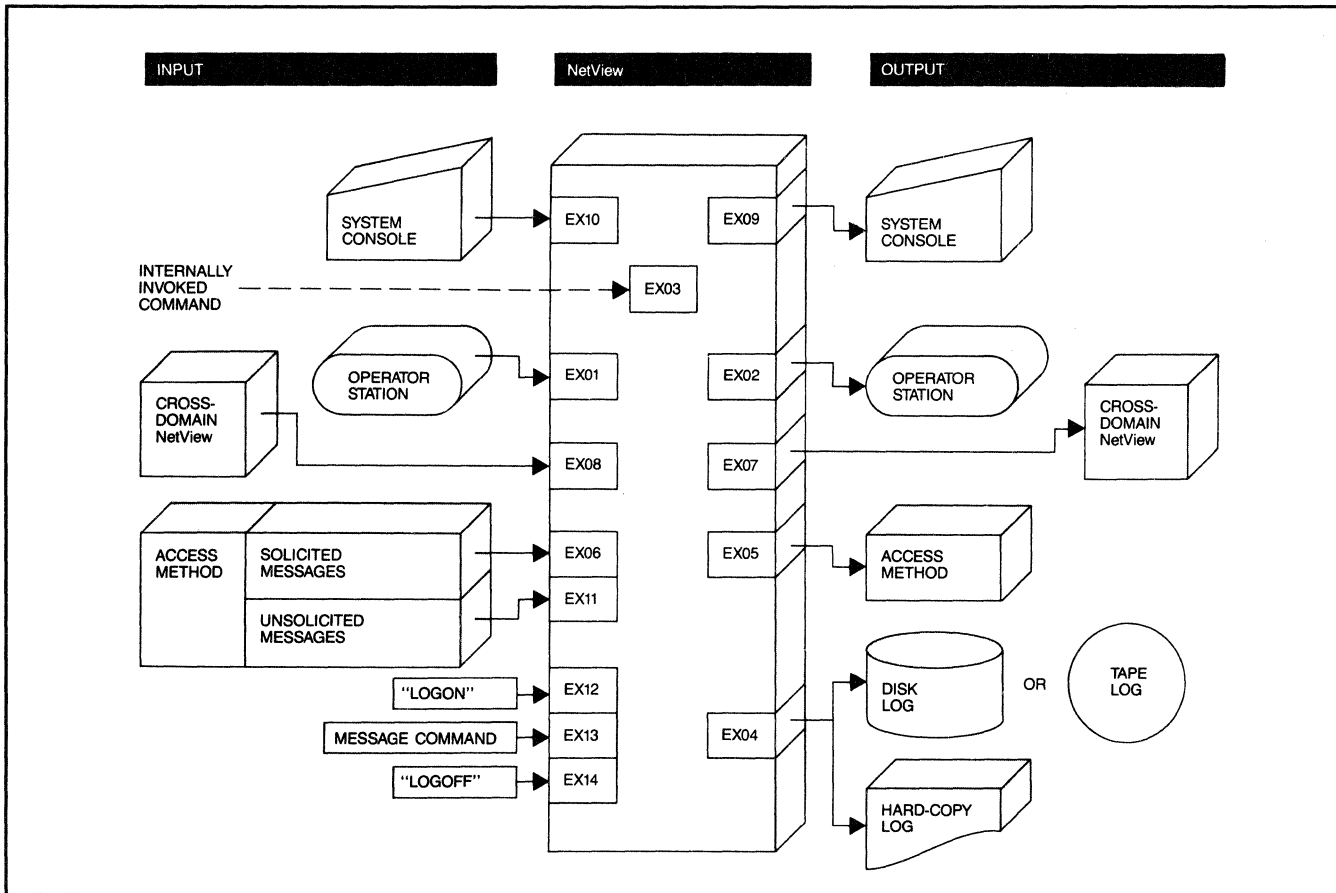


Figure 9. User exit locations.

or other attributes associated with the message. The result of these comparisons will then determine the actions taken. These actions include executing a CLIST or command processor and displaying, suppressing, or logging the message. The specifications that define the message table are in the form of IF-THEN statements and can be dynamically changed.

A sample set of message table entries and CLISTs is available with NetView to help the user get started in the automation of his or her network operations. This set automates areas of the system processes of initialization, monitoring or recovery, and shutdown for MVS, VM/SP (Virtual Machine/System Product), NetView, VTAM, JES, TSO, IMS, and CICS (Customer Information Control System). Examples are given for both single and multisystem environments. Additional customization can then be added to this base depending on the user's needs and procedures.⁸

Problem Management. The processes that take place across the problem management discipline include problem detection, the collection and analysis of data, and recovery. NetView plays an important part in

each of these processes, both directly in the function it provides and indirectly in its support for user implementation in this area.

Alerts are the main mechanism for problem detection in NetView. By monitoring the Hardware Monitor Alerts-Dynamic panel, the network operator will be kept up to date on error conditions in his or her span of control. This screen is automatically updated whenever an alert occurs, presenting a one-line summary of this special event. The operator may then select a specific alert and retrieve detailed information.

Data about potential error conditions are collected by both the Hardware Monitor and the Session Monitor components. When an event or an alert is received from a network resource, the data are stored by the Hardware Monitor in the database. These data describe conditions in the physical network. Data pertaining to the logical network are gathered by the Session Monitor. For each session, information such as names, configuration, and type is kept. If an error occurs at session activation, information including BIND and UNBIND failures and session setup fail-

IBM's Approach to Network Management

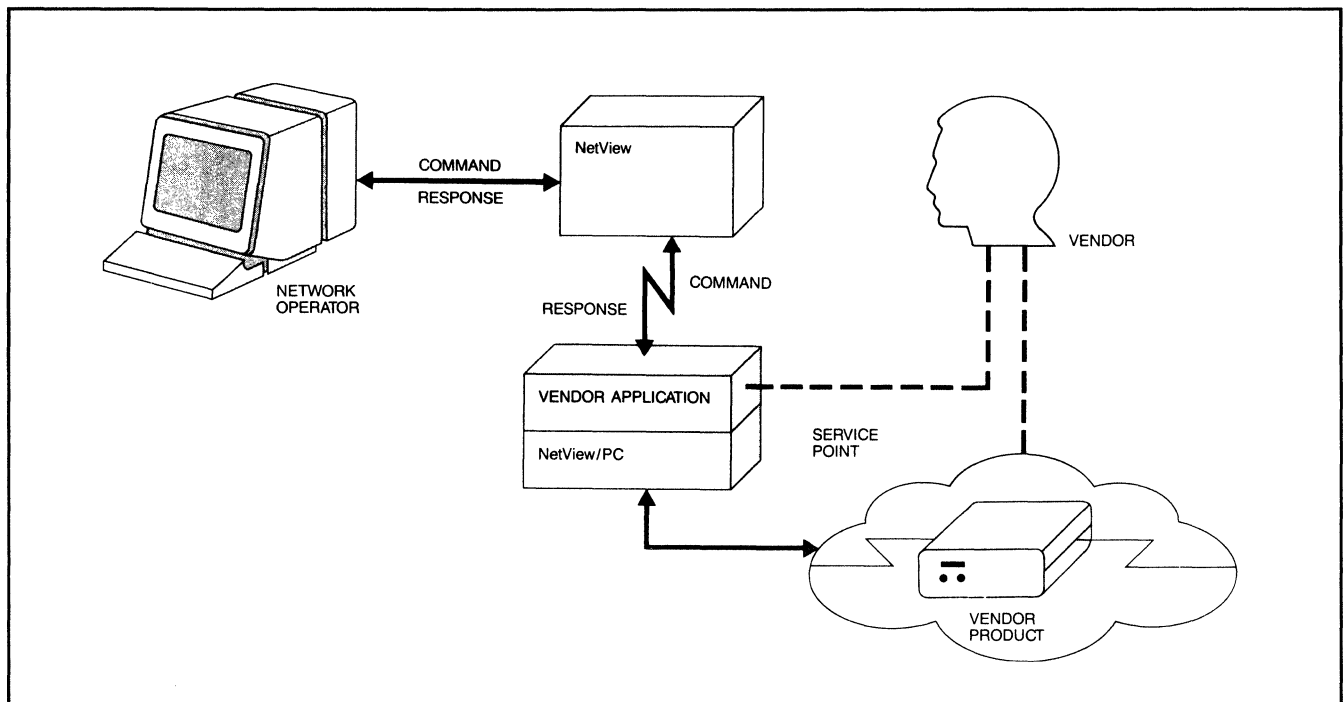


Figure 10. Service point command interface.

ures is logged. With the trace function active, PIU and NCP data are also stored on a session basis.

As part of the alert, probable cause and recommended action information is included for each situation. This analysis is done either by the sending component or by the Hardware Monitor. The goal is to isolate the network problem to the specific failing component. When statistical records are sent to the Hardware Monitor, an analysis is also done to determine whether the situation warrants an alert. On the basis of the results of these analyses, the operator may need to continue the problem determination process. Through NetView, commands can be sent to modems and NCPs to issue tests and retrieve status information.

Once the problem has been identified, the last process, recovery, can begin. In many cases, CLISTS driven from messages or alerts will automatically bypass or recover a lost resource. When operator intervention is needed, the commands can be entered from a central site to correct problems throughout the network.

Performance Management. The NetView program provides function within the performance management area that presents the user with information on response time and availability of network resources. This information can be used for problem isolation and capacity planning.

The Session Monitor collects response time data from control units having the Response Time Monitor feature. The data are associated with sessions by specification of a performance class and response time objectives. These data are displayed for the operator in summary format, by session and collection period (Figure 8), and by response time trends.

Network availability data re-collected by the Session Monitor and written to an external log for off-line processing. This information includes the time and date of the BIND and UNBIND, number of PIUs sent and received, and names of the primary and secondary LUs. The Status Monitor component provides an on-line indication of network availability. For a given set of resources, information showing the percentage of time each resource has been in each given status (Active, Inactive, Pending) is displayed.

Open Network Management Facilities

One key aspect of our network management system is that it is open, allowing participation of non-IBM devices and systems. A published set of network management architectures is provided, along with the application programming interfaces of NetView. The following subsections will explain these facilities as they relate to the focal point and to the control points.

IBM's Approach to Network Management

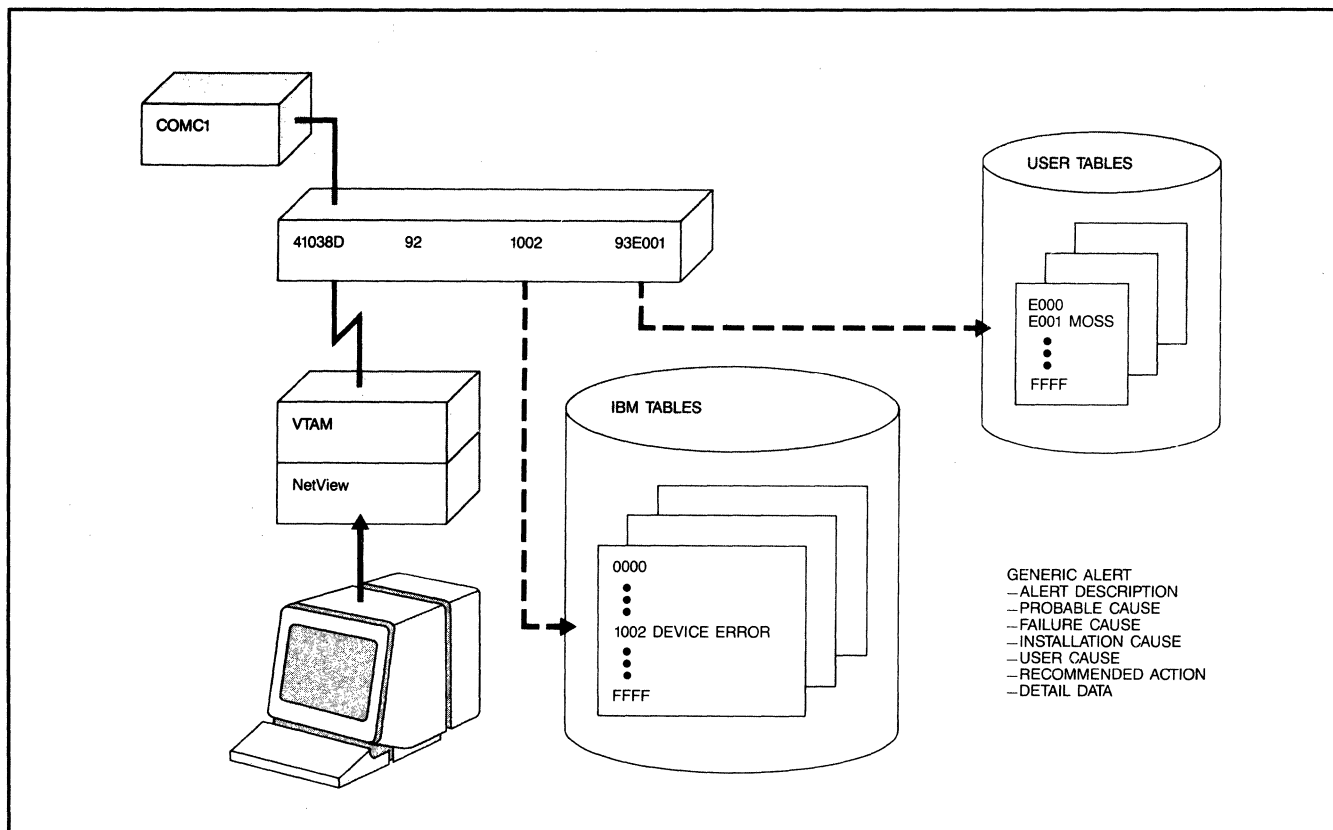


Figure 11. Generic alerts.

Focal Point Facilities. NetView contains an extensive set of application programming interfaces at the focal point for customization. These include user exits, command processors, subtasks, CLISTs, and screen modification procedures. To provide the appropriate environment for user-written programs, several NetView macros and control blocks are documented for external use. Each interface has its set of coding conventions, including input and output specifications. The facilities offered provide the user with a range in flexibility from simple panel changes to complete network management components.

Exit routines can be written to view, delete, or replace data flowing to, from, or through NetView. Exits can handle specific events, or can automate processes based on message information. Fourteen global exits in NetView apply to all tasks (Figure 9). They are located at all major data routing points. For example, data passing to and from the operator console, to the log, or to and from the access method can be changed or removed.

When an additional service or function is needed, a command processor can be written by the user. Command processors have a broader influence than exit routines. These modules are invoked by the operator in the same way as NetView commands, accepting

parameters if needed. They can output data to the operator's screen using line-by-line or full-screen mode.

If even more control is needed, a mechanism to include a user subtask is available in the NetView program. This capability is designed to allow control of a resource that is used by more than one task.

A CLIST is a set of commands and special instructions that are grouped together under one name. Commands that can be issued from a NetView terminal can be put in a CLIST, CLISTs can be invoked by an operator, another CLIST, the operator log-on procedure, NetView initialization, a user-written command processor, or a message.

The CLIST language provides comments, control and assignment statements, labels, and built-in functions. Parameters may be input with the CLIST, and a global as well as local variable capability exists. Simple control statements such as IF and GOTO are present, along with a more complex WAIT facility. It is also possible to output text to the operator's screen.

The Help facility in NetView is structured to allow the user to perform modifications and customization. The existing panels may be changed to reflect individ-

IBM's Approach to Network Management

ual procedures and practices. One or more panels may also be added or replaced at any point in the display hierarchy to allow additional help information.

Interfaces to Distributed Points of Control. In order to control, monitor, and perform problem determination on distributed resources, NetView offers two structured interfaces: Service Point Command Interface (SPCI) and generic alerts. Use of these facilities allows non-IBM or non-SNA devices to communicate with the focal point.

With SPCI, a global command, RUNCMD, is provided that allows native service point commands to be executed at the service point from the host. The service point may then in turn send a reply message to the focal point and have it displayed on the operator's screen (Figure 10). RUNCMD is sent to the service point in the Network Management Vector Transport (NMVT) format. This support⁹ allows NetView/PC to attach to NetView and provide its interface to non-SNA products.¹

Any resource in the network can send an error indication to the focal point. This indication is a generic alert. A generic alert allows coded alert data to be transported to a focal point where the data will be stored and displayed. The coded data are used as an index to a predefined table containing short units of text, or the data may contain the text directly. A user may supply his or her own table for indexing or use the one provided by IBM (Figure 11).

Generic alerts are used to produce output for the operator containing recommended actions, probable cause, and detail displays. With this facility, a customer can use the Hardware Monitor component to coordinate alert conditions from every resource in his or her network, be it IBM or non-IBM.

SUMMARY

With its comprehensive set of integrated components, the NetView program provides a strategic base for

IBM's network management system. It facilitates operational control and problem and performance management of network resources from a central point or several distributed locations. To provide openness, it offers a rich set of application programming interfaces and published network management architectures.

Network management is one of the most important needs in today's networks, and requirements in this area are growing steadily. NetView offers the facilities to meet these needs and the framework to grow along with the networks it supports. The automation support offered today allows for future development of expert-system-based diagnosis and repair of network problems. Performance monitoring, automatic load balancing, and central control and management of large multifocal-point networks are just some of the challenges to be faced and answered by the NetView program as it grows and evolves.

REFERENCES

¹For more information, see M. Ahmadi, J.H. Chou, and G. Gafka, "NetView/PC," *IBM Systems Journal* 27, No. 1, 32-44 (1988, this issue).

²*Learning About NetView: Network Concepts*, SKT-0292, IBM Corporation; available through IBM branch offices.

³*Communication Network Management Product Integration White Paper*, edited by Michael L. Nault, SHARE, Inc., Chicago (February 25, 1985).

⁴For a discussion of the general approach taken in the usability tests, see L.C. Percival and S.K. Johnson, "Network management software usability test design and implementation," *IBM Systems Journal* 25, No. 1, 92-104 (1986).

⁵*NetView Operation*, SC30-3364, IBM Corporation; available through IBM branch offices.

⁶*NetView Operations Scenarios*, SC30-3376, IBM Corporation; available through IBM branch offices.

⁷*NetView Installation and Administration Guide*, SC30-3360, IBM Corporation; available through IBM branch offices.

⁸For more information, see *Automated Operations Using NetView CLISTs*, SC30-3477, IBM Corporation; available through IBM branch offices.

⁹NetView/PC refers to this support as Service Point Command Service. □

Managing Change in SNA Networks

This report will help you to:

- Analyze the change management capabilities of IBM's SNA/Management Services (SNA/MS), as implemented in NetView Distribution Manager (NDM) Release 2, and the IBM 3174 Control Unit with Central Site Change Management facilities.
 - Implement an effective change management system to plan, schedule, and track changes to remote SNA nodes.
-

Systems Network Architecture/Management Services (SNA/MS) has been enhanced to give network users change management capabilities. The first IBM products implementing change management are NetView Distribution Manager R2 and the 3174 Control Unit with the Central Site Change Management microcode function. This report describes the design selected and the functions provided: Retrieve, Send, Delete, Install, Send-and-Install, Remove, Accept, and Activate. It also describes how SNA/MS makes use of another new SNA component designed for it—SNA/File Services,⁷ described in another paper in this issue. (Although not strictly necessary, it is recommended that the reference be read prior to reading this report.) SNA/File Services, in turn, uses an enhanced SNA/Distribution Services format to provide an architecture for file distribution in an SNA network.

Since its introduction in 1974, IBM's Systems Network Architecture (SNA) has gained wide market acceptance.^{1,2} The appearance of ever-larger SNA networks has created the need for centralized management capabilities, including the ability to assign

This Datapro report is taken from "Managing Changes in SNA Networks" by C.P. Ballard, L. Farfara, and B.J. Heldke, IBM Corporation, from the *IBM Systems Journal*, Volume 28, Issue Number 2, 1989. Copyright © 1989 International Business Machines Corporation. Reprinted with permission.

NetView, Application System/400, and AS/400 are trademarks and Personal System/2 is a registered trademark of International Business Machines Corporation.

management focal points (by geographic region, for example) and the ability to handle interconnected but independently administered networks. Increasingly diverse technologies and a multiplicity of vendors in the networking scheme have led naturally to the development of a network management component of SNA to provide these capabilities—SNA/Management Services (SNA/MS).³⁻⁵

An early requirement of managers responsible for operation of an SNA network was to be able to retrieve files containing both executable objects (for example, programs or panels) and associated data sets from one or more SNA nodes at which the files could be prepared, bring them to a central administrative site, and subsequently distribute them to remote SNA nodes for execution or processing. Also required was the ability to delete files at the remote nodes. Users needed to plan, schedule, and track all of these activities from a central site. The Distributed Systems Executive (DSX) program was developed to meet these requirements for a System/370 central site host and was first released in 1978. Over time, it was enhanced to accommodate a variety of remote systems, such as the IBM 8100, Series/1, System/36, and System/370 with the VSE operating system. More recent versions added support for IBM 4680 Store System Processors and IBM Personal Computers.

By the mid-1980s, the need to combine several network management products into one overall network management product strategy was recognized. This

Managing Change in SNA Networks




Figure 1. A network perspective.

strategy led IBM to group a number of licensed programs for network management into the NetView product family with the aim of facilitating the integration of functions and user interfaces. In October 1987, IBM announced the NetView Distribution Manager⁶ product that is based on DSX extensions and improvements. It adds support for IBM Application System/400 (AS/400), System/370 nodes with the Virtual Machine/System Product (VM/SP) or Distributed Processing Programming Executive (DPPX) operating systems, Personal System/2, and System/88.

As the variety of remote SNA nodes supported by the NetView Distribution Manager increased, requirements for architecture to support its functions became apparent. These requirements were in three major areas:

- Change management

- System-independent file management and distribution
- Bulk data transport (including store-and-forward and fan-out features)

In the case of change management, the architecture was required to support the wide variety of products participating in an SNA network and to be open so that users and other vendors could implement it if desired. SNA/MS enhancements were developed to provide the change management functions, and this report explores the architectural solution for these functions in more detail. SNA/MS makes use of a new SNA component, SNA/File Services (SNA/FS), that was developed to address the file distribution functions. Another IBM report⁷ and a reference manual⁸ describe SNA/FS in detail. The existing SNA/Distribution Services (SNA/DS),^{9,10} with streamlined formats, was

Managing Change in SNA Networks

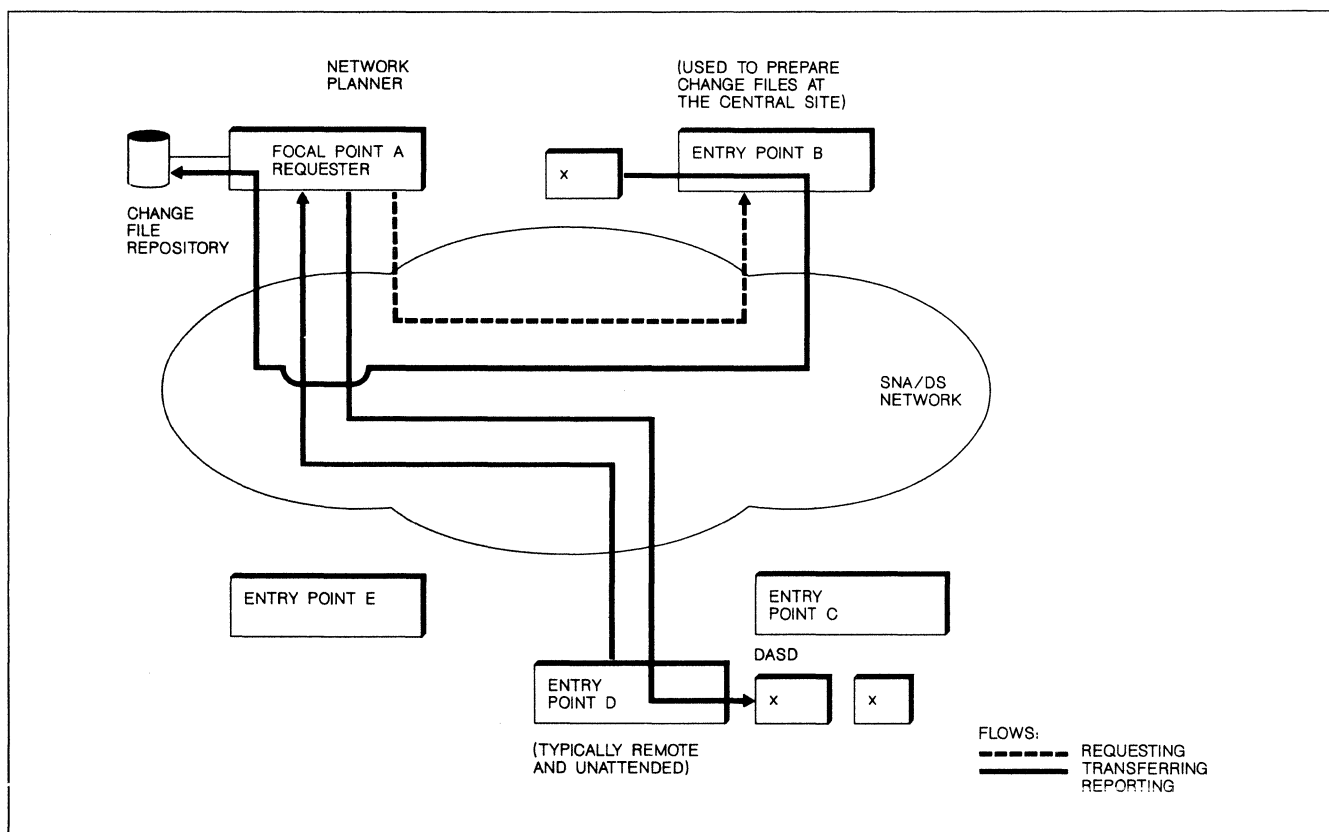


Figure 2. Architectural model for change management.

chosen by SNA/FS to satisfy the requirements for bulk data transport. The Logical Unit Type 6.2 (LU 6.2) protocol is used to exchange data between nodes participating in the SNA/DS network.

The first implementing products are NetView Distribution Manager R2 and the IBM 3174 Control Unit with the Central Site Change Management microcode function.

CHANGE MANAGEMENT ROLES AND RESPONSIBILITIES

SNA/MS defines an *entry point* as an SNA node that sends network management data (in SNA/MS format) about itself and the resources it controls to a focal point for centralized processing, and that receives and executes focal-point-initiated commands to manage and control its resources. A *focal point* is an entry point that accepts specific management and control requests of a particular network management category (for example, change management) from a user of SNA/MS (arbitrarily termed *the network planner* for architectural modeling convenience) and issues corresponding commands to other entry points.

The user has the opportunity to centralize management and control at one or more focal points.

SNA/FS capability is symmetric with respect to focal point and entry point roles. That is, nodes in either role can send files to, and retrieve files from, other nodes. Change management commands, however, are only issued from a focal point. Figure 1 shows a network of nodes participating in change management.

Entry points are typically remote and unattended. An entry point may be a large system or a small one, an intelligent workstation, a fixed function device, or a control unit. Nodes between the focal point and the entry point may perform the SNA/DS intermediate role to provide a connectionless delivery service and fan-out (that is, one copy coming into the intermediate node can be replicated and forwarded to several subsequent nodes for ultimate routing to the destinations).⁹ One of the entry points may serve as the preparation site for *change files*, that is, files containing component replacements or updates and any necessary instructions to install them. This entry point, if required, is typically located at the central site with the focal point, where it can be attended. The preparation site can also be at a focal point rather than at an entry point, or even in a system that serves neither a focal point nor an entry point role (in that case, the prepared change files

Managing Change in SNA Networks

At 14:00 EDT on 7/9/89, do the following:

```
Retrieve change file 'x' from B *
Retrieve change file 'y' from B *
Send 'x' to D *
Send And Install 'y' with 'x' as corequisite on trial at D *
Activate D using trial and production changes
```

At 14:00 EDT on 7/23/89, do the following:

```
Retrieve change file 'w' from B *
Retrieve change file 'z' from B *
Send 'x' to C and E *
Send And Install 'w' with 'x' as corequisite in production at C *
Send And Install 'z' with 'x' as corequisite in production at E *
Activate C and E using trial and production changes
```

*If this step is successful, then perform the request following

Figure 3. The network planner submits a plan to the focal point.

must be introduced by some other means, such as distribution tapes, at either a focal point or an entry point). It is the responsibility of the preparer of the change file to include in it any necessary prerequisite information to be checked by the target entry points when the change is installed. However, corequisites (a group of change files to be installed together) may be specified by the network planner.

Change Management Requests

An SNA/MS change management focal point provides for the following requests at its interface with the network planner:

- *Retrieve* obtains a change file prepared at an entry point or at another focal point for storage at the focal point.
- *Send* distributes a change file from the focal point to one or more entry points or other focal points.
- *Delete* deletes a change file at one or more entry points.
- *Install* uses a change file and its corequisites, if any, to alter, at one or more entry points, all components necessary to effect the change. The entry point can perform such alteration in a removable manner if requested, that is, so that a subsequent request (*Remove*) can return all those components to their condition prior to the alteration. The network planner can request testing either before or after the installation process is performed by the entry point if the entry point supports such testing. For example, an entry point can test a new version of a configuration file for validity before deleting the old version.

Also, automatic removal of changes (if the tests or installation fail) or automatic acceptance (see *Accept*, below) is possible. The network planner can designate components altered by the installation process for *trial* activation, or alternately, *production* activation. The designation conditions how the entry point is later reactivated.

- *Send-and-Install* is the same as *Install*, except that the focal point sends a change file in the same request.
- *Remove* returns all components previously altered in connection with a change to their condition prior to the installation of the change. It is possible only for changes installed previously in a removable manner.
- *Accept* relinquishes resources at an entry point required to maintain removability of a change and cancels the removability of a change installed previously in a removable manner.
- *Activate* causes reactivation of the entry point. Such reactivation uses changes installed on a trial basis as well as those installed in production. Also, the network planner can request that the entry point not attempt reactivation if user sessions are currently active at or through the entry point.

Testing

Without the explicit testing features provided on the *Install* request, the only testing would be done by an entry point user(s) over some period of time after the change was installed. If problems were encountered, the central-site network planner would need to be consulted, because only he or she would know about the

Managing Change in SNA Networks




Figure 4. The focal point retrieves a change file.

previous installation of changes and be in a position to issue the Remove commands.

Two kinds of explicitly requested tests are useful as part of the installation process, so that the central-site network planner can be informed immediately about the results of certain diagnostic tests as part of the installation report:

1. *Pretests*—Tests made before the programming components are altered
2. *Posttests*—Tests made after the programming components are altered but before the installation report is made

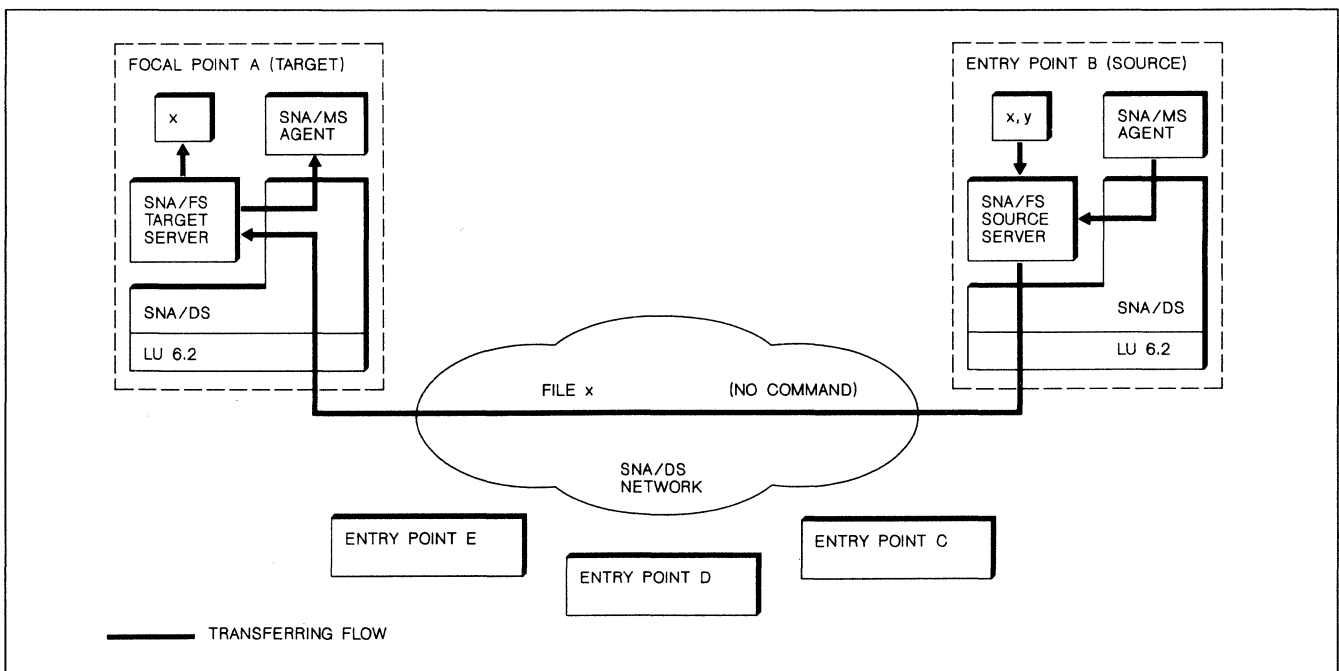


Figure 5. The change file is returned.

Managing Change in SNA Networks

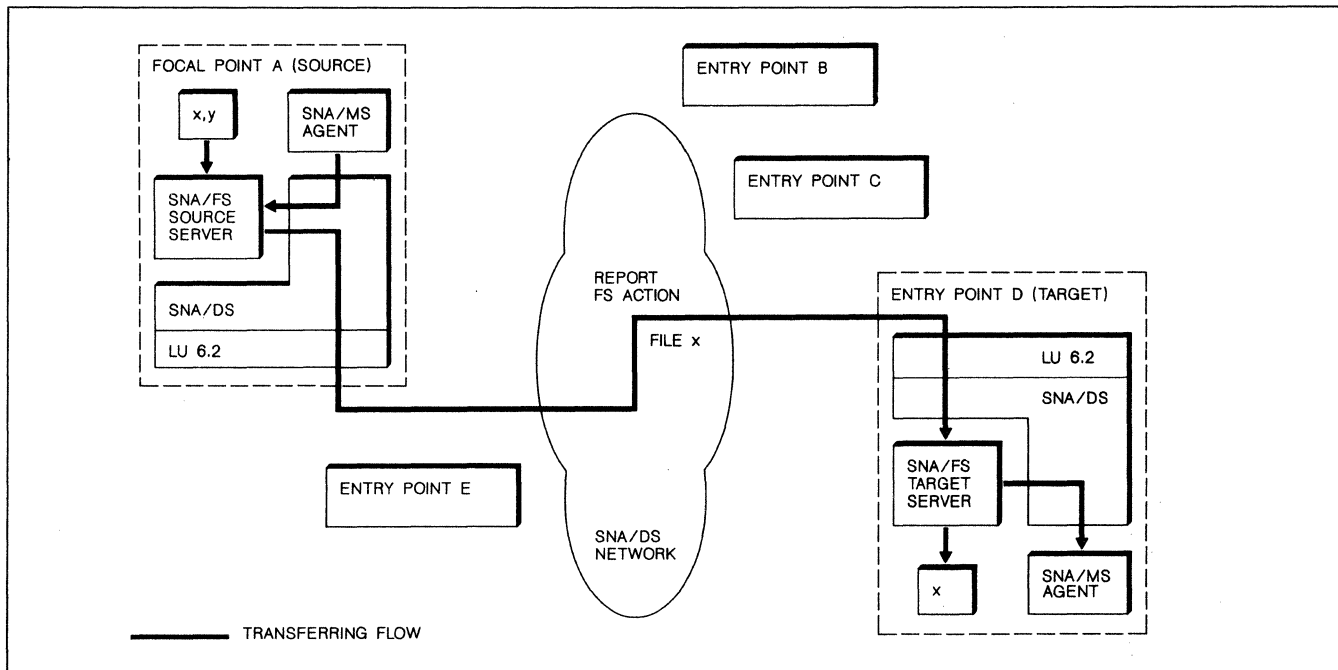


Figure 6. The focal point sends a change file.

Some reasons why automatic testing of changes is desirable:

- Possible corruption of the change between its initial development and installation at the entry point
- Possible sensitivity of the change to differences between the maintenance level of the target system and the maintenance level of those systems in which the change has already been tested
- Possible sensitivity of the change to differences between the configuration of the target system and the configuration of those systems in which the change has already been tested
- Possible sensitivity of the change to differences between functions or applications present at the target system and those present at systems in which the change has already been tested

The entry point performs the pretest (if requested) by examining the change file *before* components are altered. The change file contents may be examined for self-consistency and consistency with the configuration, maintenance level, or application set of the entry point. For example, an altered version of a file that contains input data to a routine can be checked to see if it contains inconsistent specifications.

If the pretest fails, no attempt is made to alter the components, and the installation report is made to the requester with the test results.

The posttest is performed (if requested and supported) by the entry point *after* components are altered but before the installation report is made. Altered versions of the components are tested directly, for example, by executing diagnostic routines. Such routines or test instructions can be distributed along with the change.

If the posttest fails, the components are returned to their unaltered condition if another parameter, *automatic removal*, is specified. In any event, the installation report is sent with the test results, and the requester is informed immediately. If required, the requester has the opportunity to remove the change so that the impact on end users is minimized.

The Activation Use Parameter

The *activation use* parameter of the Install request causes the entry point to install the indicated changes for *trial activation* or *production activation*. The Activate request contains a parameter that causes the entry point to activate both the trial and the production versions of altered entry point components. Entry points implement the following types of local reactivation:

1. Use of both trial and production components
2. Use of production components only

Changes that cannot be tested fully or that have a strong potential to affect the entry-point-to-focal-point communication path are best installed on trial. After

Managing Change in SNA Networks

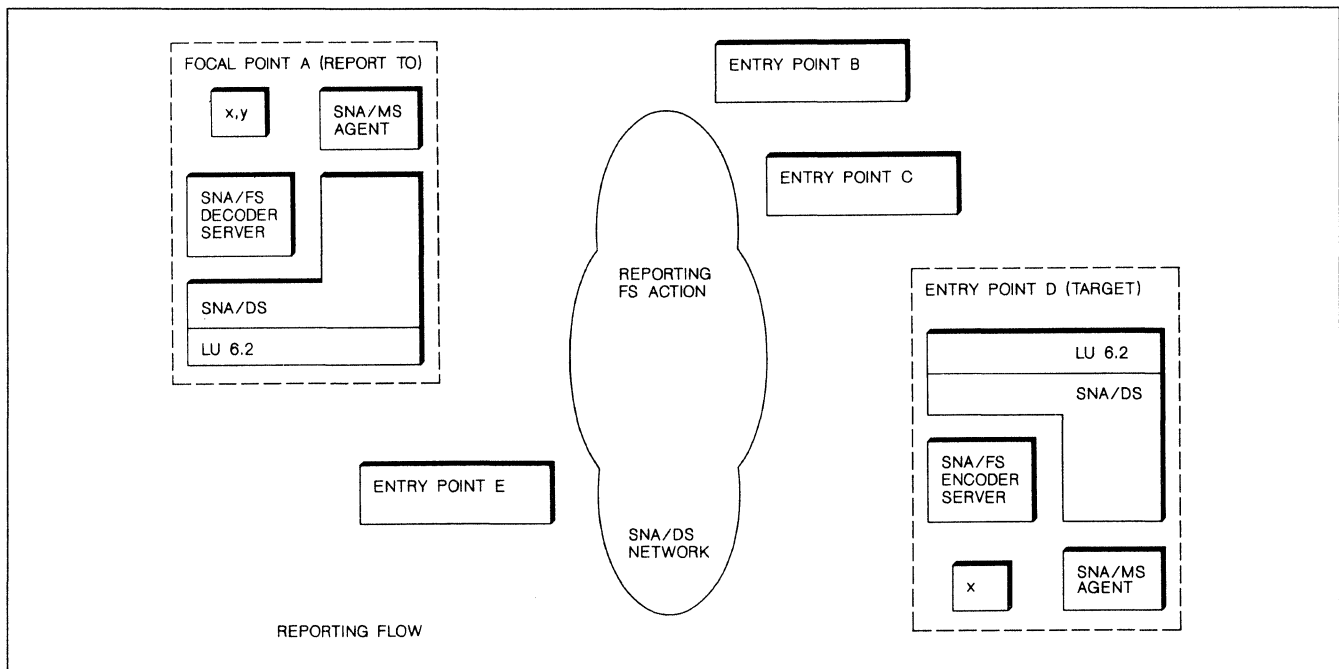


Figure 7. The entry point reports successful file transfer.

activation and a period of successful operation, the changes may be installed again in production.

Of course, to provide the trial activation capability, an entry point must be capable of installing a change so that it is removable. That is, the entry point must be able to keep a copy of the production-level system. Storage capabilities at the entry points may preclude support of trial activation in some cases. If so, the installation will be refused if activation use on a trial basis is specified.

The parameters on the Install and Activate requests reflect very specific reliability requirements of entry point implementations. Although the basics of the Activate request provide the ability to reactivate an entry point after changes have been installed (for example, by loading altered microcode), there is a need to provide a way to use only unaltered versions of components during local activation. Without such a capability, reactivation of the entry point could result in the use of a change so destructive that the path to the focal point cannot be maintained.¹¹ Repair (in the form of further change distribution and installation) cannot be triggered by the focal point in this case. Reactivation must be triggered at the entry point through human intervention. Reactivation of the (working) production-level components can restore communication with the focal point and allow the network planner to repair the components.

Hence, the ability to store both a production and a trial set of entry point components was seen as critical to

allow operator intervention to be simple, and to avoid the requirement for both change management skills and awareness of network planning activities at the entry point. Through support of a default local activation of the production system, an entry point implementation can provide hardware externals that are very simple for nontechnical users of entry points, typically not attended by technical people.

HOW THE NETWORK PLANNER USES NETVIEW DISTRIBUTION MANAGER

Focal point implementations may provide the network planner with the ability to aggregate a series of requests and specific scheduling and conditioning rules for their execution into a *change distribution plan*. In the NetView Distribution Manager implementation, plans are defined and maintained in a library, and when specified by the schedule, they are submitted for execution. The status of the submitted plan is updated as execution progresses. Control of the distribution, including recovery and restart operations, is automatically and continuously performed. The network planner builds a plan out of *phases*, each of which targets one entry point or a group of entry points. Also, files can be grouped and handled together. Each phase is built from a sequence of requests as described above. These building blocks allow the network planner to use the NetView Distribution Manager to perform the following functions:

- Start a phase at a specified date and time

Managing Change in SNA Networks

- Cancel a phase not started or completed within a specified time interval
- Join phases of the same plan with conditioning rules
- Execute a procedure at the focal point when a phase ends
- Execute a plan at the same time every day, with tolerance for a specified delay

Thus, the network planner has a powerful, easy-to-use tool to manage distribution of changes in the network, or indeed the transmission of any files. The complexity of the process is reduced because NetView Distribution Manager presents these functions to the network planner in a logical sequence, either in an interactive fashion using panels or with a batch interface.

Distribution management includes the following sub-tasks:

- Prepare and submit change management plans
- Maintain the files stored in the focal point repository
- Control the progress of submitted plans and evaluate the resulting reports
- Track the status of the change files by entry point

Multiple network planners may concurrently access the focal point facilities and perform similar or com-

plementary functions. Planning activities may be concurrent with distribution activities. Plans may be validated for correctness before their submission. Status of plan execution may be tracked by the network planner. Certain recovery actions are automatically attempted if required. Distribution may be performed unattended; however, an interactive facility is available for the network planner to monitor and control the distribution operations and take appropriate corrective actions when problems are found.

CONVERTING PLAN REQUESTS TO COMMANDS AND SERVER INSTRUCTIONS

Both focal points and entry points must contain SNA/DS support, including a specific SNA/FS server that is able to interpret control information. The application (or *agent*) using SNA/DS in this case is provided by SNA/MS. As can be seen in Figure 2, the focal point supports requests and replies at its interface with the network planner (that can be either a person or a program). A network planner request is converted by the focal point into a *command* that it sends to one or more entry points. For example, the Retrieve request is converted to an SNA/FS *Transfer To Requester* command. A command is executed by an entry point and results in a *report* to the focal point. Each report associated with a request is given to the requesting network planner when it is received by the focal point. Reports are defined by each of the SNA components. An SNA/DS report is received if an error occurred in

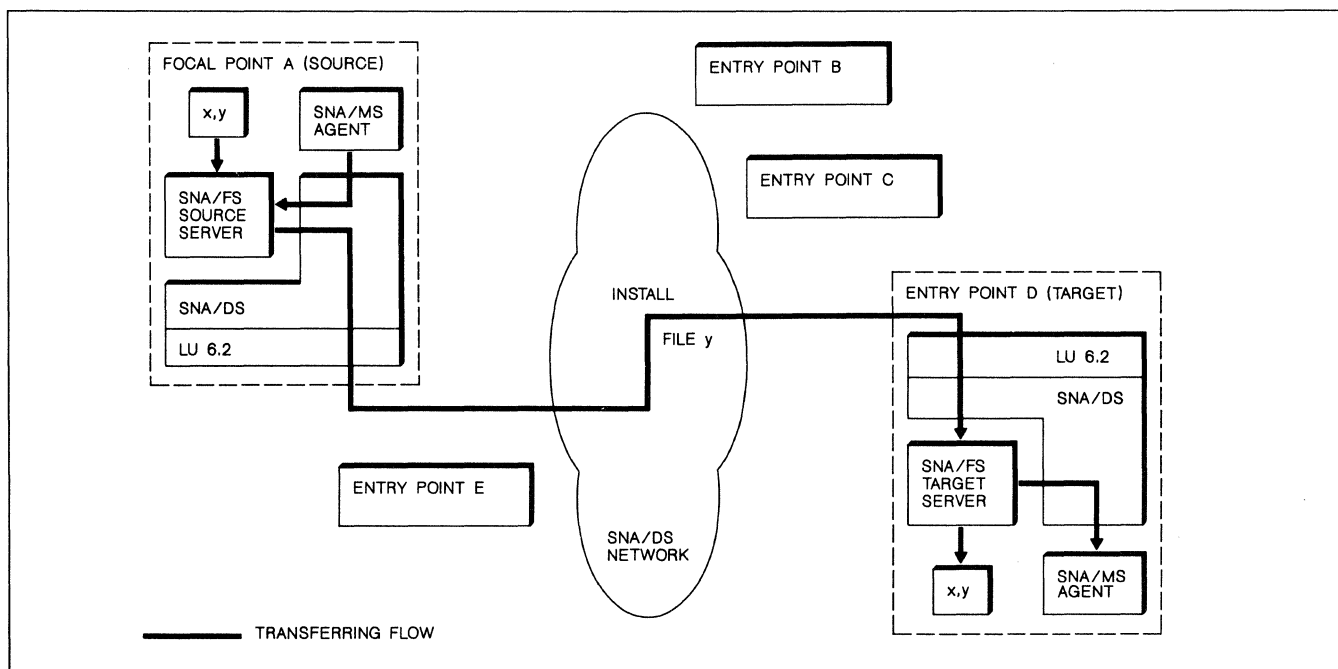


Figure 8. The focal point sends and installs a change file.

Managing Change in SNA Networks

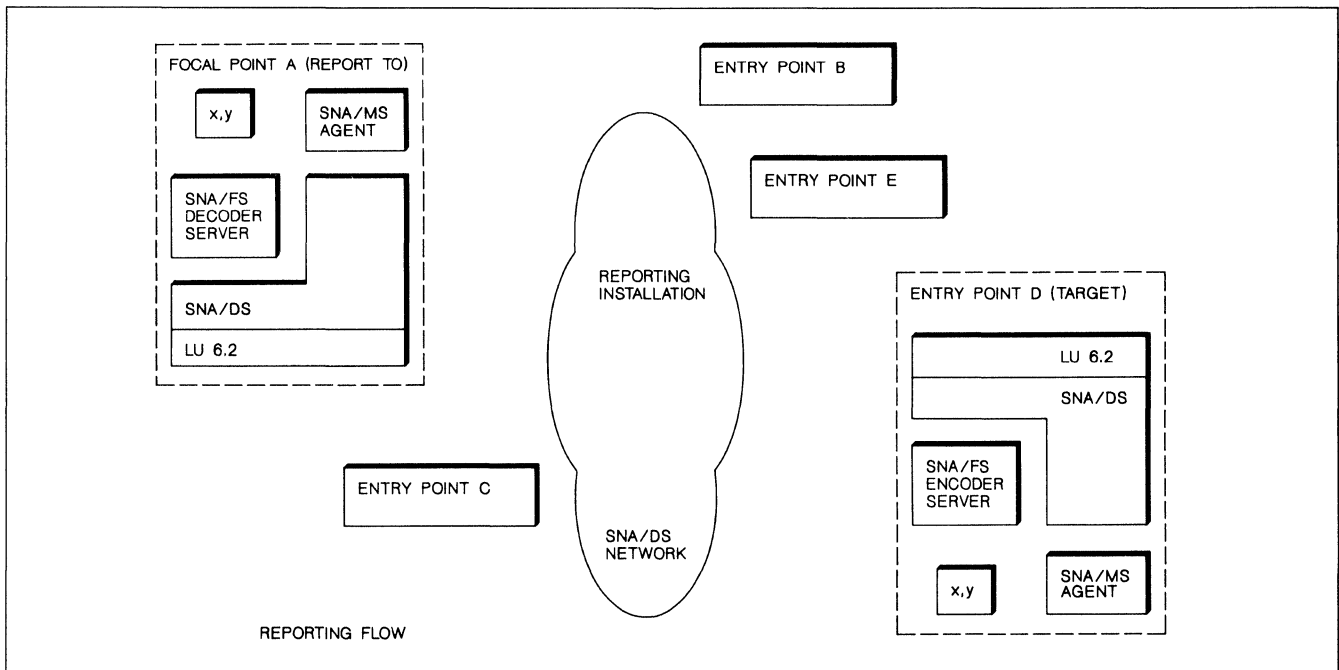


Figure 9. The entry point reports successful installation.

the distribution network. If distribution is successful, either an SNA/FS report or an SNA/MS report is received. SNA/FS reports occur if the request was for file transfer only. SNA/MS reports indicate the success or failure of the request.

Commands and reports are carried in SNA/DS *message units*. Each message unit may contain the following items:

- A command or report, defined by either SNA/MS or SNA/FS
- SNA/FS control information
- A file

A message unit may identify but not carry a file, in which case SNA/FS control information is present without a file. An example is a Retrieve request. In contrast, a command or report need not be present. For example, a successful Retrieve request results in the return of a file and its control information but no report. These examples are illustrated in the scenario given later.

Different from both network planner requests and focal point commands is the SNA/FS *server instruction* in the control information. The server instruction indicates to the SNA/FS server how to manipulate the file into, or out of, the storage facilities. For example, a Transfer To Requester Command flows with a *Fetch*

server instruction, and the reply contains a *Create&Load Or Replace* server instruction.

HOW SNA/MANAGEMENT SERVICES USES SNA/FILE SERVICES GLOBAL NAMES

One of the important motivations for this architecture is the requirement to minimize the implementation cost incurred by a focal point in identifying the files used by the wide variety of products in an SNA network. SNA/FS provides a global name for a file, consisting of a set of tokens. For a description of this feature, please see Ashfield and Cybrynski.⁷ SNA/FS standardizes the number of tokens allowed (up to 10), the size of each token (up to 16 characters), total name length ($65-n$ where n is the number of tokens), and the character set allowed for the token values (a limited set of character graphics displayable on most types of displays). In addition, SNA/FS architecture maintains a registration of values for the highest order token. For example, MCODE is the registered value for change files containing microcode. SNA/MS maintains a registration of the values of some of the other tokens, delegating authority in some cases to administrative organizations. For example, the IBM machine type is the second token for microcode. As a result, a focal point is required to implement only one input panel, say, for a user to identify microcode files for a potentially wide variety of types of target entry points.

Managing Change in SNA Networks

Definitions of token values by SNA/MS were made to satisfy two general requirements: First, each file containing a change must be uniquely identified; second, the user must be provided with some idea of the type and identity of the change when displaying the name, or an application must be allowed to process the token values. For example, the file name MCODE.9135.NA.PATCH.1234 indicates that the file contains a patch rather than an engineering change. This is needed to uniquely identify the file and also to provide useful information on display.

It is advantageous for the user to create (or have provided by product developers) change files that can be installed on a large number of entry points. Such change files allow the fan-out feature of SNA/DS to be fully exploited and reduce the user's effort. For example, files containing microcode can be designed to be applicable to many control units, whereas those containing customizing data are specific to individual control units.

The types of change files applicable to SNA/MS are microcode, customizing data, software, procedures, applications data, and documentation. The first product offering and architecture release support microcode and customizing data. Within each of these types, SNA/MS defines some of the lower order tokens, and some additional token values are defined by product implementations.

SNA/MS makes use of SNA/FS *partial name processing* for both retrieval and distribution. Partial name

processing is used when the network planner wishes to specify only some of the file identification tokens (typically, the higher order ones). An example of this is when the latest version of a customizing data file is to be retrieved, but the user cannot remember the version number contained in the lowest identification token (and does not care what its value is).

On distribution, partial name processing may be needed to specify which change file to destroy to make room for a new one when entry point storage constraints arise. An important requirement addressed by the architecture is that destruction take place only when installation has been requested properly and pretests have been performed successfully. The complete identity of the file replaced is included in the report of successful installation.

A SCENARIO

The following scenario illustrates how the network planner uses the change management functions provided by SNA/MS.

Consider a network planner working at a focal point A, whose job is to prepare changes at a local entry point B and distribute them to a number of remote entry points—C, D, and E. A new microcode change 'x' is received from IBM and prepared at B for network installation. Since a corresponding change to customizing data is required to use a new function introduced by 'x', the network planner also prepares three change

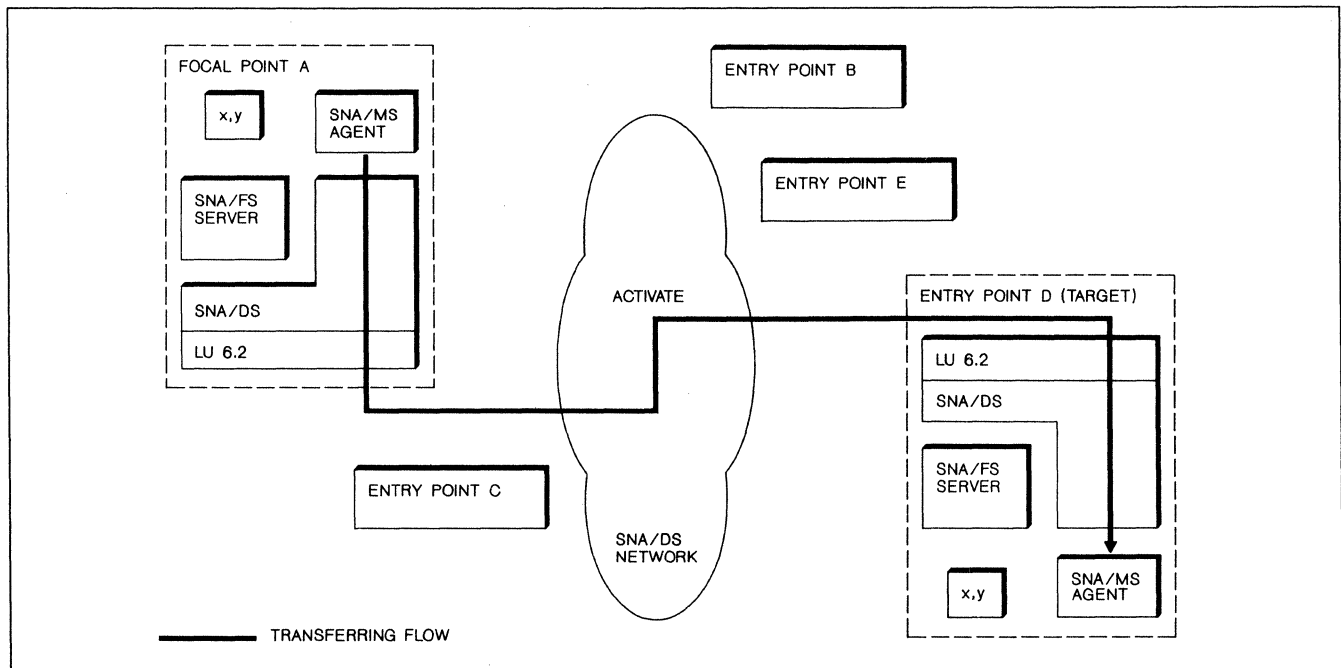


Figure 10. The focal point reactivates an entry point.

Managing Change in SNA Networks

files, 'w', 'y', and 'z', containing the customizing data for entry points C, D, and E, respectively. This preparation is done at entry point B, where the files are created and stored. The network planner prepares and submits to the focal point the change distribution plan shown in Figure 3.

The following discussion shows how the plan is accomplished by the focal point according to the architecture. The accompanying figures illustrate the process.

First, the Retrieve request is converted to an SNA/FS Transfer To Requester command. The SNA/FS server at the focal point encodes a server object containing the file name, and the SNA/FS server at the entry point decodes it (Figure 4).

The SNA/MS agent at entry point B instructs its SNA/FS server (in the *source* role) to fetch the file and its SNA/DS component to send it to the focal point. The file is sent without a report, and the SNA/MS agent at the focal point is signaled by its SNA/FS server that the file has arrived (Figure 5). The SNA/FS server at the focal point (in the *target* role) stores the file into the repository, and the next request in the plan can proceed. For a description of the SNA/FS server roles, please refer to Ashfield and Cybrynski.⁷

The retrieval of the change file 'y' containing the customizing data is performed successfully in the same manner.

The next request is to send 'x' to D. The SNA/MS agent at the focal point instructs its SNA/FS server (in the *source* role) to fetch 'x' from the file repository and builds and sends an SNA/FS Report FS Action command that is carried with the file (Figure 6).

The SNA/FS server at D (in the *target* role) stores the file, and the entry point SNA/MS agent obeys the command by building and sending an SNA/FS report to the focal point (Reporting FS Action) (Figure 7).

Next, the customizing data are sent, and installation of both 'x' and 'y' as corequisites is requested in the same flow. In this case (Figure 8), the file is sent with an SNA/MS Install command (instead of an SNA/FS Report FS Action command). The SNA/FS server at the entry point stores the file, and the SNA/MS agent installs 'x' and 'y'.

An SNA/MS report (Reporting Installation) is built and returned to indicate successful installation (Figure 9).

Reactivation of entry point D is requested, and the command is sent to D (Figure 10). The SNA/FS servers at the focal point and entry point are not involved.

The entry point reports that reactivation will be attempted, and the plan is complete.

Reactivation of the entry point includes termination of the session between A and D since the entry point completely reinitializes itself.

Two weeks later, at the time and date specified in the plan, the customizing data for the two remaining entry points are retrieved and the new microcode is sent and installed at C and E as well, this time in production. If the new microcode had caused problems at D, the network planner could have canceled the plan to prevent this final step.

SUMMARY

SNA/Management Services has been enhanced to provide change management capabilities. A manager responsible for operation of an SNA network can use it to plan, schedule, and track changes to SNA nodes that are typically remote and unattended, in a nondisruptive fashion, during their normal operation. Application developers can use it to distribute and install new software and associated application data throughout a network where an application is distributed.

REFERENCES

- ¹R.J. Sundstrom and G.D. Schultz, "SNA's first six years: 1974-1980," *Fifth International Conference on Computer Communication*, Atlanta, GA, North-Holland Publishing Co., Amsterdam (September 1980), pp. 578-585.
- ²R.J. Sundstrom, J.B. Staton III, G.D. Schultz, M.L. Hess, G.A. Deaton, Jr., L.J. Cole, and R.M. Amy, "SNA: Current requirements and direction," *IBM Systems Journal* 26, No. 1, 13-36 (1987).
- ³D.B. Rose and J.E. Munn, "SNA network management directions," *IBM Systems Journal* 27, No. 1, 3-14 (1988).
- ⁴R.E. Moore, "Utilizing the SNA Alert in the management of multi-vendor networks," *IBM Systems Journal* 27, No. 1, 15-31 (1988).
- ⁵*SNA/Management Services Reference*, SC30-3346, IBM Corporation (1989); available through IBM branch offices.
- ⁶*NetView Distribution Manager General Information Manual*, GH19-6587, IBM Corporation (1989); available through IBM branch offices.
- ⁷J.C. Ashfield and D.B. Cybrynski, "System-independent file management and distribution services," *IBM Systems Journal* 28, No. 2, 241-259 (1989).
- ⁸*SNA/File Services Reference*, SC31-6807, IBM Corporation (1989); available through IBM branch offices.
- ⁹B.C. Housel and C.J. Scopinich, "SNA Distribution Services," *IBM Systems Journal* 22, No. 4, 319-343 (1983).
- ¹⁰*SNA/Distribution Services References*, SC30-3098, IBM Corporation (1989); available through IBM branch offices.
- ¹¹Testing is optional at installation time, and implementations of testing cannot always verify the change. Software or microcode defects generally cannot be detected during an Install test function. □

Network Management in APPN Networks

This report will help you to:

- Grasp Advanced Peer-to-Peer Networking (APPN) network management fundamentals.
- Configure NetView DM/DSNX to achieve change management and distribution in APPN networks.
- Configure and use Distributed Host Command Facility (DHCF) to perform fault analysis in APPN networks.

AS/400 NETWORK MANAGEMENT

IBM's Advanced Peer-to-Peer Networking (APPN) architecture is relatively new. Full implementation became available to System/36 users in 1986 and to AS/400 users in 1988. APPN differs from IBM's traditional Systems Network Architecture (SNA) primarily because midrange systems, such as the S/36 and AS/400, are much freer to access each other *directly* through the network—rather than indirectly via SNA's hierarchical framework. Mainframes may participate in APPN networks, but they are not required.

AS/400 network management functions are performed by OS/400 Communications and Systems Management (C&SM) modules and supporting microcode. There are three functional C&SM modules:

- **Distributed Host Command Facility (DHCF)**—DHCF communicates with a Host Command Facility (HCF) program to provide mainframe-based *remote operations* of AS/400 facilities.

- **Distributed Systems Node Executive (DSNX)**—DSNX communicates with either a host NetView DM 2.2 or DSX 2.3 program to achieve mainframe-based *change management* and distribution.
- **OS/400 focal point and alert services**—These services support network *problem management* within the context of an APPN network. Also, OS/400 alert services alone can support network fault management when the AS/400 is communicating with SNA hosts or Low Entry Networking (LEN) nodes running NetView products.

This report examines in detail these three facilities' relationship to NetView, describing the functions they support, and demonstrating how AS/400 implementation differs from previous network management facilities on the System/36 and System/38.

This report was developed for Datapro by Elinor Gebremedhin, an independent data communications consultant and free-lance writer. Ms. Gebremedhin possesses a wide background in both mini and mainframe systems and software; her current areas of focus include IBM and IBM-compatible systems, distributed processing, and communications software.

Index to This Report	Page
Relationship with NetView	302
(DHCF)—Remote Operations	302
DSNX—Change Management	305
Problem Management (SNA/APPN and NetView Alert Support)	307
PC Considerations	309

Network Management in APPN Networks

RELATIONSHIP WITH A NETVIEW HOST

Mainframe Network Management Products: Each IBM mainframe network management software product is a discrete entity and sold as a separately licensed program. Of these products, only those under the NetView umbrella are key to APPN network management—since the AS/400 problem management facilities and DSNX have the capability to participate in NetView-managed SNA networks. The NetView software family currently includes 9 key mainframe products and 13 PC (or PS/2)-based packages.

AS/400 Network Management Products: AS/400 HCF, DSNX, and NetView-type fault management facilities are separate logical networking modules. IBM also incorporates them as standard elements of a single OS/400 operating system licensed program package. As of June 1988, the OS/400 operating system was at modification level 1.2. This software was automatically shipped to all AS/400 users on December 9, 1988.

The OS/400 operating system includes facilities for handling both hierarchical SNA networks and midrange peer-to-peer networks. NetView is basically a hierarchical management system—its primary function is to route alerts to enable operators to manage the network from a single location. The AS/400 C&SM functions must use SNA facilities when communicating with the mainframe, but need not do so when handling nodes subordinate to itself. For example, when AS/400 DHCF or DSNX communicate with mainframe HCF or NetView DM, the AS/400 uses the SNA Upline Facility (SNUF)—which is not part of APPN. Then the AS/400 acts as an intermediate node for either of the mainframe programs. It uses APPN to route the mainframe requests to subordinate nodes.

SNA continues to evolve and grow. Consequently, IBM must adapt NetView at each significant growth step. Announcements of upgrades for many of the mainframe and PC products were made on September 20, 1988. To accommodate that growth and change, IBM also announced upgrades for many mainframe and PC network management products. On the same day, IBM extended SNA's explicitly defined pathways from 8 to 16 routes per destination and extended sub-area addressing from 256 to 65,000 nodes (thus easing problems encountered when using 9370s as intermediate systems with large numbers of attached workstations and terminal clusters).

These changes indirectly affect AS/400 communications software. By supporting SNA networks of this size and complexity, IBM is, in effect, inviting users to interconnect secondary focal points for network management. This creates areas of control which are inter-

connected via the main network backbone, similar to the way Token Ring LANs can now be connected.

IBM added major automation capabilities to NetView at the same time it adjusted NetView to support expanded SNA networks. The shaded sidebar on this page outlines the current level and major functions of each software module and the primary growth areas provided by the most recent release.

THE EVOLUTION OF C&SM FUNCTIONS

AS/400 C&SM functions are SAA-compatible facilities embedded in the OS/400 operating system. C&SM functions evolved from, and participate in, several different types of management networks originally designed for older systems.

This evolutionary development has produced a somewhat complex network management operating environment. There are different user interfaces, configuration rules, and problem analysis facilities for each of the different functional modules: DHCF (remote operation), DSNX (change management), and problem management. Thus, each module is equivalent to a separate network. Furthermore, the set of other computer systems that can interconnect to each module is not the same, and neither are the available communications protocols and network support functions.

DHCF—REMOTE OPERATIONS

The ability to configure an AS/400 for host-based remote operations evolved from existing System/36 DHCF programs. These programs, in turn, evolved from software dating back to the mid-1970s.

The AS/400 DHCF interfaces to a mainframe partner program, the Host Command Facility (HCF). This configuration reduces or eliminates the need to station data processing professionals at the remote site. Although HCF was originally intended to allow only *mainframes* to control midrange systems as dependents, AS/400s now have the capability to control sister systems as if they were subordinates. This is achieved by using 3270 emulation and routing the control logic through a mainframe host.

Mainframe users invoke the Host Command Facility (HCF) from 3270-type terminals to perform remote operations (such as running an application) on AS/400 systems. HCF and its partner DHCF manage the interaction: The AS/400 sees the terminal as if it were an attached 5250-type display; the mainframe 3270 terminal sees the AS/400 application as if it were on the

Network Management in APPN Networks

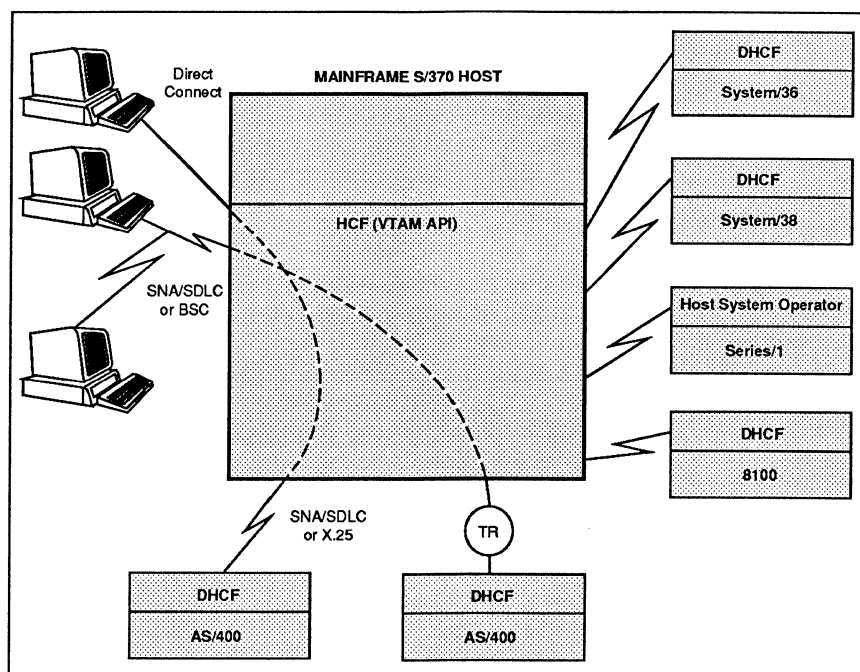


Figure 1. This figure depicts several host terminal users operating remote AS/400 systems. With the exception of a few keyboarding restrictions, this interaction is transparent once it is established.

local mainframe. (See Figure 1.) Once established, this interaction is transparent for the most part, although there are a few keyboarding restrictions due to the mapping back and forth between 3270- and 5250-type terminals.

Since the connection is transparent, the mainframe user can use the 3270 to perform any operations normally accepted by the AS/400 from a 5250-type terminal—providing both the mainframe and the AS/400 have equivalent authority. The mainframe user can access, control, and run applications; display user files and libraries; and look at and respond to unique AS/400 system messages.

The mainframe user's capabilities depend on the nature of the target DHCf system and not the connection itself. But there can be some lack of transparency in accessing "the same" (i.e., ported) application running on different systems. At present, an AS/400 application that runs in equivalent form on other HCF/DHCf system types (like the 8100, System/36, etc.) would probably be slightly different, even if the AS/400 version had been converted from the other system.

This lack of transparency is one of the problems SAA is designed to eliminate, or at least reduce. At present, impure transparencies between different DHCf systems running the same application depend largely on what the user has done on his/her own in the past to control application portability problems.

HCF/DHCf was not originally designed to allow role reversals. Previously, the distributed system could not

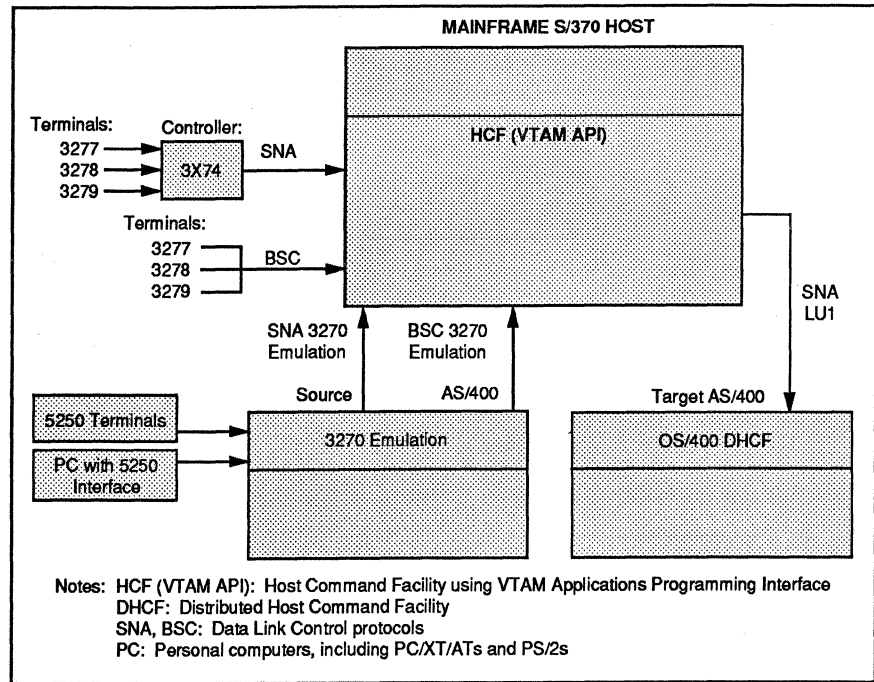
access and run mainframe programs in the same manner. Recent DHCf releases have allowed one AS/400 or System/36 system to operate another AS/400 or System/36 in the same way that a mainframe host does, providing that the interaction is routed through a mainframe host.

The source AS/400 5250-type terminal invokes 3270 emulation software (standard on the AS/400) to present its request to the mainframe. The mainframe sees the source AS/400 as if it were a 3270-type terminal; the mainframe sees the target system as if it were another 3270-type terminal. The source request is routed to the target system and converted back to 5250 formats using the 3270-to-5250 conversion software available on the AS/400 (and the System/36). The application program on the target AS/400 thus sees the incoming request as if it were made by a local 5250-type terminal. The configuration underlying this type of interaction is shown in Figure 2. For more information on source/target communication in APPN networks, see *Datapro Reports on Communications Software*, "IBM AS/400 Advanced-Level Peer-to-Peer Networking (APPN)," Report CMS20-491-601.

HCF/DHCf is a general-purpose facility and is not specifically designed for change management or problem analysis. It does allow mainframe users to access any of the AS/400 operations or service facilities to perform problem analysis, including running and displaying of storage dumps and traces, and interactive examination of the system's error log.

Network Management in APPN Networks

Figure 2. This figure illustrates the use of 5250 to 3270 emulation on source AS/400s to control target AS/400s running DHCF. While HCF/DHCF was not specifically designed for change management or problem analysis, it does allow mainframe users to access AS/400 operations and facilities to perform problem analysis tasks—such as running and displaying storage dumps and traces, and interactively examining the system's error log.



HCF/DHCF was developed much earlier than NetView, however, and therefore cannot directly support NetView's semiautomated alert handling capabilities.

Configuring DHCF

A 3270-type terminal can access an AS/400 DHCF system through an HCF host in one of two ways. It can communicate with the mainframe as a remote device via a BSC or SNA/SDLC connection. Or it can communicate with the mainframe as a local device via an SNA or non-SNA direct connection. In the second instance, the mainframe host communicates with the AS/400 using LU1 SNA/SDLC or X.25 line protocols.

As shown in Figure 1, several mainframe 3270 terminal users can simultaneously access one AS/400 DHCF system. Up to 254 terminal users can be accommodated on each communications line.

The DHCF software on the AS/400 is part of the OS/400 operating system and runs on any basic AS/400 configuration that supports 5250 terminal connections. The HCF Version 2 licensed program (5668-985) is required on the host. HCF can operate under VTAM/NCP or under the Network Communications Control Facility (NCCF) together with the Terminal Access Facility (TAF) feature and VTAM/NCP. HCF operating environments include MVS/370 Release 3.8 or later, OS/VS1 Release 7.0 or later, or DOS/VSE Release 5.0 or later.

Recent releases of HCF/DHCF support newer distributed systems (such as the AS/400), but remain compatible with older operating environments. The AS/400 requires an upgrade to HCF Version 2 on the mainframe; using Version 1 will produce unpredictable results.

Comparing AS/400 DHCF and Systems/36/38 DHCF

The key differences in AS/400, System/36, and System/38 DHCF support revolve around defining and selecting 3270-to-5250 keyboard mapping procedures. System/36 uses a KEYS procedure that is not supported on either of the other two systems. System/38 uses a Define Keyboard Mapping (DEFKBDMAP) command. The AS/400 replaces this command with a Set Keyboard Mapping (SETKBDMAP) command, and adds two others, Change Keyboard Mapping (CHGKBDMAP) and Display Keyboard Mapping (DSPKBDMAP).

DSNX—CHANGE MANAGEMENT

Mainframes can use NetView Distribution Manager (NetView DM) software to distribute system changes, programs, and data to a variety of midrange systems in an SNA network. The partner program on most of the midrange systems is the Distributed Systems Node Executive (DSNX), originally developed for the 8100. The AS/400 version is very similar to the one used for System/36 APPN networks. An AS/400 can act as an

Network Management in APPN Networks

intermediate node for distributing DSNX objects to other AS/400s, to System/36s, or to PCs attached through an SNA Distribution Services (SNADS)/APPN network.

NetView DM evolved from the Distributed Systems Executive (DSX). DSX was originally developed to distribute changes to 8100 systems running a DSX-related form of the Distributed Systems Node Executive (DSNX). The AS/400 version of DSNX, oriented toward mainframe NetView DM, is included as part of the standard OS/400 software.

With the exception of NetView itself, DSNX and NetView DM are undoubtedly the most important programs introduced for the NetView environment. DSNX and NetView DM allow mainframes to automate centralized control of distributions. This capability is essential not only to network problem management, but also to the day-to-day, run-of-the-mill requirements of distributed processing. These routine requirements include the requesting and disbursement of active database files, save files, other objects such as temporary files, batch jobs, messages, user interface support facility functions, and others.

Change management can be achieved in networks exhibiting either a two-tiered or a three-tiered structure. In two tiered-networks, NetView DM distributes all files and information to AS/400s, which act as end nodes. In three-tiered networks, AS/400 DSNX systems can act as intermediate nodes as well, by processing distribution lists received from the host and then forwarding requests and/or changes to appropriate System/36s, other AS/400 systems, and/or personal computers.

AS/400 DSNX consists of three modules: the *host interface*, the *request processor*, and the *DSNX/PC support processor*. Two other programs are also required when the AS/400 acts as an intermediate node (passing DSNX messages to other end points): the OS/400 object distribution facility and SNA Distribution Services (SNADS).

The host interface module: When the AS/400 operates solely as an endpoint, the host interface module passes NetView DM requests from System/370 hosts and routes in one of two ways: through object distribution to the DSNX request processor, or through SNADS to the DSNX/PC module. The AS/400 receives responses through object distribution or SNADS and forwards them back to the host through the SNA Upline Facility (SNUF). SNUF is started by the host and must run under a subsystem description that specifies a communications entry (which includes a default ID) and a routing entry (for the SNUF device that communicates with the host).

The DSNX request processor: After processing the host's NetView DM request, the DSNX request processor sends its responses back to the host interface module by means of the object distribution facility. This processor runs under the OS/400's QGPL/QDSNX subsystem. When the QDSNX subsystem is inactive, host requests are placed in a queue until activation.

The DSNX/PC support module: The DSNX/PC support module receives host requests through SNADS, places each request on a queue, and waits for the personal computer to request the queued host by means of a APPN/APPC (LU6.2) session. The "*APPC" default routing entry is usually used. (NOTE: Local Token Ring LANs are viewed as types of SDLC lines by the network software.) The personal computer's response is routed into SNADS for delivery to the host interface node.

Once set up, DSNX processing and distribution facilities involve very little operator interaction, even when controlling intermediate nodes are active in processing. DSNX processing and distribution facilities can be used for upgrades and changes in system-resident programs, with the exception of the operating system changes.

Configuring NetView DM/DSNX

A NetView DM/DSNX network is a hierarchical network in which a centralized host downloads programs, and controls changes, for a number of nodes. The host is a 30XX, 43XX, 9370 or other System/370 mainframe running either NetView DM 1.1, or Distributed Systems Executive (DSX) 3.2 with a Program Temporary Fix (PTF). The host program can control nearly all of IBM's major midrange systems, but curiously, the System/38 is excluded from DSNX networks, as shown in Table 1.

The OS/400 DSNX configuration affects NetView DM/DSNX communications requirements. In the basic application configuration, the OS/400 DSNX is a simple subordinate application like other midrange nodes in an SNA network. In this case, the OS/400 DSNX system is an endpoint—directly related to the host NetView DM program, but possessing no relationships with other DSNX nodes.

The OS/400 DSNX may also be configured as an intermediate node, passing DSNX objects to other AS/400s, System/36s, or (generic) PCs. (See Figure 3.)

PCs attached to AS/400s or System/36s must receive distributions through a SNADS/APPN network facility. This requirement holds even when the PCs are at-

Network Management in APPN Networks

Node Environments	NetView DM*/ DSNX	DSX 3.2/ DSNX	DHCF	APPN Alert Focal Points
AS/400 as OS/400	—	—	X	X
AS/400 as S/36 SSP	X	X	X	—
System/36 SSP	X	X	X	—
System/36 as S/1 CPS	X	—	X	—
System/38	—	—	X	—
Series/1 CPS	X	X	X	—
Series/1 EDX	X	X	X	—
Series/1 RPS	X	X	X	—
System/88	X	—	—	—
30XX/43XX/9370 VSE/SP 2.1	X	X	—	—
3174 Terminal Cluster	X	—	X	—
3790 Comm. System	X	—	—	—
4680 System	X	—	—	—
8100 DPPX	X	X	X	—
8100 DPCX	X	X	X	—
Personal computers	X	X	X	—

Note: * NetView DM mainframe host environments are MVS and VM. They do not operate as DSNX.

Table 1. AS/400 Communications and Systems Management (CS&M) network nodes.

tached to an OS/400 which communicates directly with (or, is local to) the host. In this instance, the intermediate OS/400 DSNX node passes the distribution to SNADS. SNADS “knows” that the PC is directly attached to either a communications device or to the Token Ring LAN. Figure 4 depicts this logical pathway.

Since DSNX is part of the OS/400, there are no separate system software requirements for DSNX. AS/400 DSNX communications software requirements are dependent on the configuration type:

Simple Application: This configuration requires SNA Upline Facility (SNUF). SNUF, an OS/400 operating system module, is integral to supporting communications with NetView DM. In this type of configuration, other systems in the NetView DM network, such as Series/1 or 8100, have no effect on the AS/400 configuration.

Intermediary for AS/400s and/or System/36: The configuration requires SNADS (a part of OS/400) as well as SNUF. APPN is usually used as well. In this type of configuration, only APPN hosts (excluding System/38) can be included in the network.

Intermediary for PCs: requires SNADS and APPN, both part of OS/400 as well as SNUF. PCs can include the PC, XT, AT, and PS/2 in all their permutations. Again only APPN hosts (excluding System/38) can be included in the network.

Comparing AS/400 NetView DM/DSNX and System/36 NetView DM/DSNX

DSNX networks do not include System/38. While OS/400 DSNX and System/36 DSNX function similarly for the most part, however, there are three aspects in

which OS/400 DSNX differs from the earlier System/36 version. These three aspects are file compatibility, queue management, and logon:

- **File compatibility:** With the exception of files, there is little object compatibility between System/36 DSNX and OS/400 DSNX. Files are object compatible and can be routed through a host in order to be exchanged as database members. Since the System/36 and the OS/400 have different naming conven-

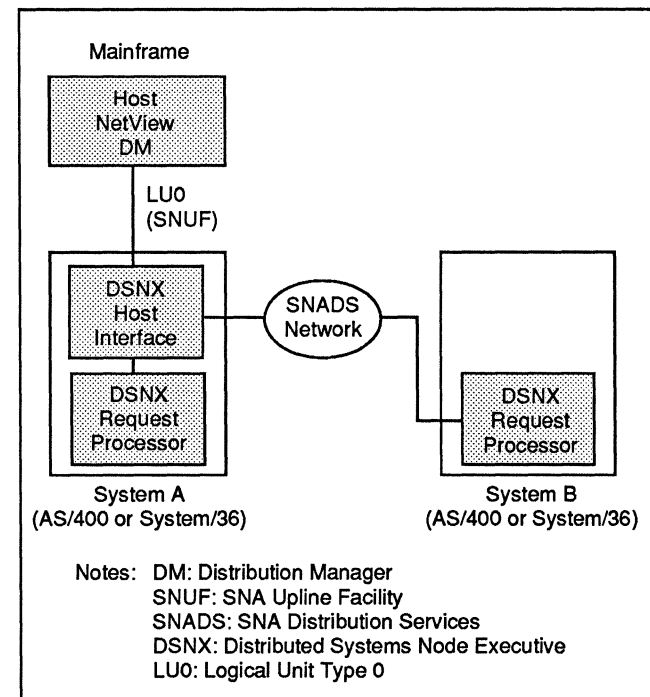


Figure 3. DSNX request processors may reside on a physically separate system or on the same systems as the DSNX host interface.

Network Management in APPN Networks

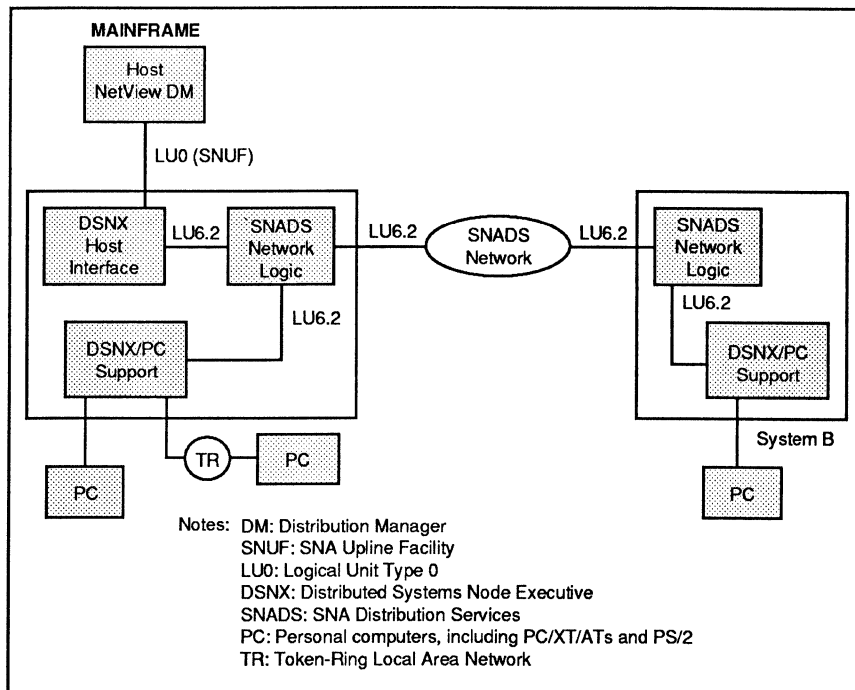


Figure 4. This diagram illustrates the logical distribution pathway between NetView DM and PCs. The PCs may be attached to a communications device or a Token Ring LAN.

tions, the host system must rename the files when handling this type of swap.

- **Queue management:** OS/400's "Work With PC/DSNX Queues" command (WRKDPCQ) only allows the user to display or delete queue entries; the equivalent functions on System/36 also allowed the user to hold or release entries or entire queues.
- **Logon:** If there is no active session in process, OS/400 automatically performs a logon to the host when replies are received from previous requests. The System/36, provides an external interface, so a user can start a session. The AS/400 does not provide this feature.

PROBLEM MANAGEMENT (NETVIEW AND ALERTS)

SNA alerts supply a network operator with information on the actual or impending loss of availability of a resource. This alert information includes all available problem analysis data concerning the event. There are several ways to handle the process of creating, sending, and logging SNA (NetView) alerts within networks that include AS/400 systems. An AS/400 can act as:

- **The primary focal point in an APPN Network.**
- A secondary APPN focal point. In this case, the AS/400 passes screened alerts up to the NetView host.

- **An intermediate or end node (EN) in an APPN or SNA network.** In this case, the AS/400 passes alerts without acting as a focal point.

The critical element in configuring this type of management network is the "focal point" computer system that collects and analyzes the alerts for those systems within its "sphere of control."

For more information on focal points in SNA network management, see "IBM SNA and NetView," Report NM40-491-101.

NetView was originally structured to manage networks from a single mainframe "primary focal point." IBM has since evolved NetView to include secondary or "nested" focal points, allowing users to develop a three-tiered system. The bottom tier is always a "Service Point Command Facility" (SPCF), usually a PC, that can forward alerts and act on commands received from the focal points. SPCFs cannot serve as focal points, however.

SNA Alert Support

Alerts are messages sent to the local system operator or a remote operator. Alerts contain information about problems with software errors or hardware resources, including local devices or their controllers, communication lines, or remote devices or their controllers. AS/400 alert support includes creating alerts, sending and receiving alerts, logging alerts, holding alerts, and displaying alerts. The AS/400 can also act as a problem

Network Management in APPN Networks

management focal point, or provide "sphere of control support." (See the section entitled "Management Services Sessions" for information on sphere of control.) Since the AS/400 can act as a problem management focal point, the AS/400 has the option of receiving and displaying alerts, and taking appropriate action directly instead of forwarding the alerts to mainframe NetView.

Unlike System/36 APPN Network Nodes, AS/400 APPN Network Nodes can act as primary focal points—directly receiving all "alerts" concerning existing or potential problems. AS/400 APPN Network Nodes can also act as lower level focal points that receive and forward alerts directly to the primary focal point (usually a mainframe host), with or without being screened. This type of three-tiered, nested-focal point design can do much to reduce network traffic and mainframe loading.

Management Services Sessions

Management services sessions define the nature of the focal point and its allowable interactions. The focal point's "sphere of control" is the collection of network nodes and end points from which it receives alerts. A system is *not* part of the sphere of control just because it is connected to a primary focal point. Rather, a system must run management services sessions according to an explicitly defined sphere of control. Management services sessions maintain connectivity with other network nodes, accept alerts received from systems in the sphere of control, and forward alerts if a higher level focal point exists.

A default focal point can also be defined. A default focal point receives alerts from any system lacking an assigned primary focal point. The default focal point asserts the relationship status to the new system and begins receiving alerts without any additional actions on the user's part. The sphere of control can be automatically expanded to include the new systems.

The AS/400's APPN-oriented network management services sessions are unique at present: System/36, System/38, and System/370 cannot participate in interactions between these types of network nodes. The AS/400 may also conduct a more limited type of network management using "alert controller" sessions. The AS/400's network attribute is defined as an alert controller session under any one of the following three conditions:

- The AS/400 is sending alerts to the System/36, System/38, or a System/370
- The AS/400 is not using an APPN network

- The AS/400, for any reason, is not using management session services

The user must define a controller description for sending alerts in order to establish alert controller sessions. Consequently, the responsibility for handling the alerts is left up to the receiving system. IBM advises that management services sessions should be used instead of alert controller sessions, if at all possible.

Figure 5 depicts the variety of connections supported by an AS/400 primary focal point. The figure shows three types of connections to System A: The first con-

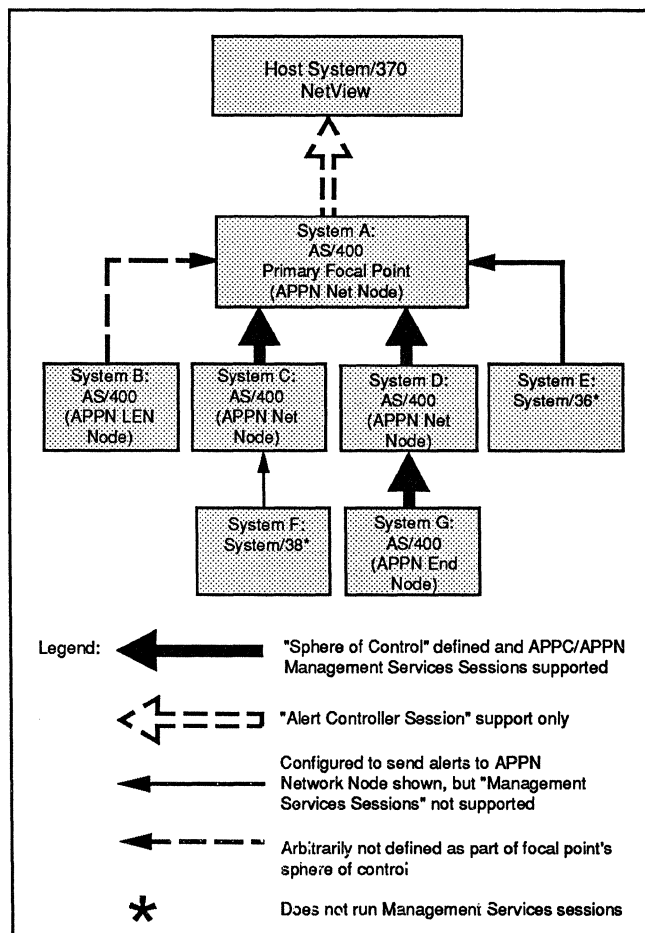


Figure 5. This figure depicts the variety of connections supported by an AS/400 primary focal point (System A). First, System A is connected to two lower level secondary focal point systems (C and D) which are part of A's sphere of control. Second, System A is also connected to a System/36 (E) which forwards NetView alerts but is not part of System A's sphere of control. Third, System A is connected to System B which, although part of the APPN network, does not forward alerts or participate in the NetView network at all. In addition, A's sphere of control extends to APPN End Node System G. System A's sphere of control does not extend to System F since, as a System/38, it lacks the requisite Management Services Sessions software.

Network Management in APPN Networks

nection is from System A (the AS/400 primary focal point) to lower level secondary focal point systems (C and D) which are part of A's sphere of control. The second connection is from System A to System E, which forwards NetView alerts but is not part the sphere of control defined by System A's management services sessions. The third connection is from System A to System B, which arbitrarily is part of the APPN network but does not forward alerts or participate in the NetView network at all. In addition, the sphere of control extends to an APPN End Node System G which has been able to participate in System A's sphere of control . . . unlike System F which does not have the requisite Management Services Sessions software to do so.

Differences in AS/400 and Systems/36/38 Alert Support

There are two primary differences between AS/400 and System/36/38 alert support. First, the AS/400 is designed to use APPC/APPN and management services sessions to interact with other systems that support similar capabilities. Second, the AS/400 is designed to define which systems are within the sphere of control of a focal point. System/36 and System/38 do not have these capabilities. AS/400 system connections to systems that lack management services sessions must be handled by the ALeRt ConTRoL Unit attribute (ALRCTLU) of the CHAnGe NETwork Attributes (CHGNETA) command. System/38 uses a similar ALRCTL parameter on CHGNETA, capable of handling SSCP-PU sessions, but System/36 uses an altogether different CNFIFICF procedure to configure an APPC/APPN subsystem that can handle both SSCP-PU or PU-PU sessions with other systems.

On both the AS/400 and System/38, alert generation is controlled using the alert status (ALRTSTS) network attribute. On System/36, the ENABLE procedure command is used to enable the appropriate APPC or APPN subsystem—alert generation is then initiated using an ALERT procedure in conjunction with a pre-defined subset of system messages.

Operator-generated alerts are sent to the AS/400 and System/38 with the Analyze Problem (ANZPRB) command and to System/36 using the ALERT NOTIFY procedure. On all three systems, messages CP19804, CP19805 or CP19806, designed for general use, can also be used for this purpose.

Alerts are logged into the QALERT journal in the QUSRSYS library on System/38 when the alert focal point (ALRFOCPNT) network attribute is “*YES”. On the AS/400, received alerts or locally generated alerts are controlled by the alert logging status (ALR-

LOGSTS) attribute; alerts are logged into a physical file, QALERT in library QUSRSYS. On System/36, alerts are logged only when they cannot be sent; in that case, they are logged in the ALERTFIL disk file.

Alert primary focal point (ALRPRIFP) and alert default focal point (ALDFTFP) network attributes on the AS/400 are not the same as the System/38 ALRFOCPNT network attribute.

PC CONSIDERATIONS

Personal computers, including PCs, XTs, ATs, PS/2s and their various permutations, are connected into NetView networks by means of the multitasking NetView/PC program. NetView/PC is a very important program, because it operates as the Service Point for the Communication System (SPCS) not only for PCs, but also for other devices as well.

Key PC Functions

Key PC functions include an Alert Manager, a Problem Manager, a Service Reminder, DOS Partition support, Data Logging and Report Generation Module, Remote Console Support, the Applications Programming Interface/Communications Services (API/CS), the Communications Manager, and a General Help Function.

The Alert Manager either automatically sends alerts to a connected host, or stores alert data for later display and handling by a local operator. Personal computers running NetView/PC cannot operate as secondary focal points.

API/CS is an interface between device-dependent applications written in PC Macro Assembler and NetView/PC. This open interface allows a network operator and network management applications to access key services that support problem management—either locally or via transfers to mainframe NetView management facilities. IBM itself supplies several PC programs using this interface, most notably the interface to the Token Ring.

Among services available through this interface are file transfers to CICS/DDM, alert transfers to NetView and exchange of C & SM data with user-written applications.

The Communications Manager is based on APPC/ LU6.2 facilities, a PU Type 2 identification to host ACF/VTAM and NCP, and either asynchronous or SDLC synchronous line connections. Switched synchronous or asynchronous line speeds must be 2400

Network Management in APPN Networks

NETVIEW COMPONENTS AND OPTIONS

NetView Release 3: Provides the base for handling all network alerts in VM/SP, VM/XA, MVS/370, and MVS/ESA environments. NetView at the Release 2 level is also available for DOS/VSE. The original NetView, introduced in 1986, included facilities for problem management, automatic operations, performance monitoring and accounting facilities. Release 3, available May 1989 for MVS/XA/ESA and August 1989 for VM, allows a focal point to monitor alerts within its own or any interconnected network domains, supporting the distinction between primary and secondary focal points within a network. This release also adds support for PL/1, C, and Knowledge-Tool interfaces; adds IBM LAN Manager support; and allows NetView command lists to be written in the high-level, SAA-compatible REXX language.

NetView Distribution Manager (NetView DM) Release 2: Provides transfer of special-purpose files and other objects to support change and problem management between MVS and VM hosts and their usually subordinate DSNX nodes. DSNX nodes can include AS/400, System/36, Series/1, 8100, System/88, 4680, 3174, VM and VSE systems (e.g. 9370) as end nodes, but three-tier networks can only use AS/400 and System/36 as intermediate nodes, and AS/400, System/36 and personal computers as end nodes. NetView DM was originally introduced in 1988 for MVS and became available to VM users in January 1989. Release 2, available June 1989 for MVS/SP and December 1989 for VM/SP, makes use of facilities provided by September 20, 1988 enhancements to ACF/VTAM Version 3 Release 2 and later releases, SNA Distribution Services (SNADS), a new SNA File Services (SNA/FS) component, a new change management category of SNA Management Services (SNA/MS), added programming interfaces, and extensions to the General Interactive Executive (GIX) and the Batch Utility (BU). All of these components, together with the new release, are required to support direct communications among multiple change management focal points, and to allow central site control of 3174 microcode and configuration data.

NetView Performance Monitor Release 3: Originally released in 1987, NetView Performance Monitor significantly extends the network performance tuning and accounting capabilities included in the basic NetView program. IBM is continually adding new features, such as Network Gateway Accounting support (available March 31, 1989).

NetView File Transfer Version 2: Supports generalized transfers of VSAM KSDS, QSAM, single PDS members, and physical sequential data sets on DASD or tape. Version 2 provides for transfers between MVS, VM and VSE environments in peer-to-peer node relationships.

NetView Network Definer Release 2: Reduces the effort and skill needed to create and update definition tables for VM/SP-based SNA networks. Release 1 was first available November 1988. Release 2, available in March 1990, adds a full screen interface, as well as interfaces to IBM Token Ring LANs, leased SDLC lines, and X.25 permanent virtual circuits.

NetView/Access Services 1.1.1: Enables a user on one MVS or VM terminal to gain single access control to a number of different applications concurrently, while protecting the network applications against unauthorized use. This program was originally introduced in 1987. Version 1 Release 1.1 for MVS, available May 1989, also provides an interface to the SNA Application Monitor (SAMON).

Information/Management (Info/Man): Provides three databases: Problem Management, Change Management, and Configuration. The Problem Management database details the reporting, tracking, and closing of problems. The Change Management database tracks the planning, coordinating, and monitoring of changes. The Configuration database keeps an inventory of hardware and software components in the network.

SNA Application Monitor (SAMON) Version 1 Release 3: Allows MVS terminals to display status information about any of the applications running in a network as well as about the network itself, and to gain access to any of its applications. Release 3, available May 1989, adds adaptations to MVS 2.2, refresh of SAMON 1.2, extended security processing, and PF-key assignments for applications functions and commands.

SolutionPac NetView for Automated Operations: Packages Release 2 of MVS NetView with other NetView programs to provide a cohesive, tailored set of SNA host network management services.

bps, but nonswitched synchronous lines can be 9600 bps for long haul lines, or 2400, 4800, or 9600 bps for digital lines.

NetView/PC Programs

The major separately licensed programs related to NetView/PC are as follows.

Network Management in APPN Networks

NetView/PC Version 1.2: Release 1.1 allows PCs, XTs, ATs and PS/2s running under PC-DOS 3.2 or 3.3 to automatically send alerts to NetView mainframes or AS/400 primary or secondary nodes; the Realtime Interface Co-Processor (RIC) is also required. Main memory taken up by DOS, RIC, and NetView ranges from 330K to 460K bytes out of the total addressable 640K bytes. NetView/PC 1.2, available April 28, 1989, extends NetView capability to PS/2 systems operating under OS/2, enhances generic alerts, allows up to 128 applications to connect to a chainable SPCF, and provides for asynchronous communications access by means of an API/CS.

PC Node Executive (PCNX) Version 1: Allows NetView DM changes to be extended to XTs, ATs, and PS/2s attached to MVS/XA or MVS/ESA hosts. Version 1, available March 31, 1989, allows the PC to initiate transmissions from the focal point at defined times and dates, as well as to execute changes in an unattended mode.

VTAM Protocol Conversion Application (VPCA) Release 1: Allows a PC or PS/2 running PCNX to communicate with NetView DM by converting LU Type 0 protocols to LU Type 2 and vice versa, and managing the conversation. Release 1, available March 1989, also performs error recovery and security/authorization checking for the PCNX sessions initiation.

LAN Manager Entry Version 1.0: Provides management of single-segment Token Ring LANs or broadband or baseband PC Network LANs from a single OS/2 EE 1.1 station. This program, available March 24, 1989, can be an upgrade to the PC 3270 Emulation LAN Management Version 1.0 introduced in 1987 (usually used with NetView/PC.) The LAN Manager Entry Version 1.0 problem management facilities include critical resource monitoring and an application alert transport service.

LAN Manager Version 2.0: Provides management of single or multiple segment Token Ring LANs with/without bridged broadband PC Network LANs from a single station. This station can be a local LAN Manager station or a remote NetView console. Version 2.0, available March 24, 1989, runs under OS/2 EE 1.1, and can use its communications facilities or can coexist with NetView/PC 1.2 on the same station. Using CLIST commands, the host NetView console operator can access 11 LAN management functions concerning adapter, bridge, and network status functions.

PC Network Bridge Version 1.0: Bridges two Token Ring LAN segments, two broadband PC Network segments, or one Token Ring segment connected to a broadband PC segment. The product supports either

4M bps or 16M bps Token Ring segments. Release 1, available May 1989, is designed to alleviate LAN limitations on numbers of devices attachable, host connections, and overall network expansion capability. The product detects problems on supported network segments and forwards the alerts to the NetView LAN Manager.

Token Ring Network Bridge Version 2.1: Bridges two Token Ring LAN segments, including both 16M bps and 4M bps (in any combination). Release 2.0 was available January 27, 1989; Version 2.1, available July 1989, adds remote bridging and filtering of forwarded frames for remote bridges.

Like the PC Network Bridge, the Token Ring Bridge programs support problem detection of the bridged network segments. Once problems are detected, the product forwards alerts to the NetView LAN Manager.

NetView Voice Network Design Version 1.0: Examines the cost-effectiveness of installed routing configurations in voice networks, selects the minimum cost mixture of offnet facilities, consolidates summary files, and provides a tariff database. The program, first released on March 31, 1989, can use Network Call Accounting or customer-supplied call accounting detail records.

NetView Network Call Accounting Version 1.0: Helps the customer control operating costs of telecommunications assets, including calls placed from PBX and CBX or Centrex systems. This product, first delivered March 31, 1989, is a focal point voice management application.

Matrix Switch Host Facility (MSHF) Release 2 (5688-0091). MSHF allows the NetView console to control the IBM 3728 Matrix Switch. MSHF2, available October 28, 1988 in the U.S. and November 1988 in Canada, extends this capability to control of the larger AUTOSWITCH 1000 and 4000 series products.

Transmission Network Manager (TNM): Provides a PS/2-based control point for Integrated Digital Network Exchange (IDNX) networks. Release 1 can operate in standalone mode or within a NetView network using included NetView/PC facilities.

Enhanced Operator Console: This program, originally treated as an IDNX hardware feature, has been available through IBM Software Distribution as of September 20, 1988. The Release 1 licensed program version is an easy-to-use AT or PS/2-based system capable of MIS report generation.

Alert Monitor: This program, originally treated as an IDNX hardware feature, has been available through

Network Management in APPN Networks

IBM Software Distribution as of September 20, 1988. The Release 1 licensed program version is a NetView/PC-based system capable of real-time access and/or automatic retrieval of IDNX events and alarms, automatically reformatting them to be handled as SNA-defined alerts.

All of the above programs were originally designed to forward alerts to a local operator or to a mainframe host for handling—neither AS/400 nor System/36 alert processing was incorporated into the products' architectural design.

As stated earlier, AS/400 change management is a new capability for NetView, and it is not available to System/36 or System/38. Furthermore, the AS/400 cannot simply define a PC as within its sphere of control; the PC must be configured to send its alerts to the AS/400.

Similarly, the NetView/PC version of DPNX is mainframe-oriented. Hence, DPNX requests and responses are routed through APPN and SNADS to the AS/400, and then routed to a mainframe—even if the PC is locally attached through the workstation interface or a Token Ring LAN. (AS/400 communications software treats Token Ring LANs as a type of SDLC communication line.) The PC is not treated as if it is directly attached.

ADDITIONAL OBSERVATIONS

In HCF networks, one AS/400 can also operate other AS/400s, although this capability is outside of APPN. As explained earlier, the source AS/400 system emulates the usual local terminal, so the host is not really operating like an intermediate node in a three-tiered network. HCF systems are wholly linked by SNA 3270 emulation sessions.

In NetView networks, it is only when APPN/APPC is used that the AS/400 can become a primary or secondary focal point managing a sphere of control. In either capacity, the AS/400 may or may not be forwarding alerts to mainframes or other AS/400 systems. If APPN is not used, only alert controller sessions (that forward the alerts to another system) can be used. Neither System/36 nor System/38 is capable of the change management sessions that allow the AS/400 to be defined as a focal point, so the introduction of the AS/400 version of APPN has added new network management capabilities in the midrange. Mainframe-based SNA and APPN have developed new relationships.

NetView DM supports attachment of a number of different midrange systems as end points, but only AS/400s or System/36s operating in the context of APPN networks can serve as intermediate nodes that pass requests on to other AS/400s, System/36s or personal computers. □

T1 Network Management: A Strategic Perspective

This report will help you to:

- Evaluate network management capabilities of T1 multiplexers.
 - Select the best T1 networking equipment based on specific management criteria.
 - Perform alarm, circuit, and inventory management for T1 networks.
-
-

Today's T1 multiplexers have brought an unprecedented degree of flexibility and functionality to T1 network operators. It is exactly this functionality and flexibility that make increasingly sophisticated network management systems a necessity in today's networking environment.

Early T1 networking equipment had little need for effective network management as we know it today. First-generation multiplexers were essentially point-to-point channel banks that did not have any real-time reconfiguration capabilities. Network managers could insert a card at each end of a circuit and, depending on the kind of cards, establish a circuit essentially on a hard-wired basis. Any configuration of ports required site personnel to set switches on individual cards.

Second-generation muxes added drop-and-insert capabilities that allowed circuits to be established directly from one port to another on different multiplexers. These too were fixed connections that featured, at the very most, static routing tables—a feature that provided some flexibility in configuration circuits but demanded a great deal of operator intervention for each circuit or change desired.

Third-generation multitasking multiplexers allowed network managers to define alternative routes for a circuit to take if service on the primary link was interrupted. Drop-and-insert or bypass capabilities were added that allowed a circuit breaker between two end multiplexers to be routed through an intermediate multiplexer. It was the task of the operator, however, to describe exactly what steps the network should take in the event of facility failures—and to manually update the routing tables.

Today's software-based fourth-generation muxes provide a host of routing and configuration alternatives. All changes can be performed in software by a single operator. In addition, intelligent routing algorithms automatically select the optimal routing path for circuits. The network manager simply tells the system which two endpoints to connect, and the multiplexers will automatically build the routes. In the most sophisticated systems, network nodes are continually evaluating conditions in the network, so that circuits will be optimally routed based on the configuration of the network at the time a route is requested. (Figure 1 provides diagrams of these four types of T1 multiplexers.)

EVALUATING THE OPERATOR INTERFACE

The two common ways to approach the interface between the network and the operator are text-based and graphics-based systems. Most systems today use a text-based interface to manage all network func-

This Datapro report is based on "T1 Network Management: A Strategic Perspective," by T. Musselman, Digital Communications Associates, Inc., from *Telecommunications*, February 1988. © 1988, Horizon House-Microwave, Inc. Reprinted with permission.

T1 Network Management: A Strategic Perspective

tions. Operators can search a specific node for information about the operating conditions of that node. The text-based interface also allows the operator to “zoom in” to view specific circuit detail and examine operating parameters for the ports at each end without the need to make any additional menu selections. One virtue of a text-based interface is that it allows operators to access network management functions from any location by dialing in from an ASCII terminal (e.g., a VT100-type terminal) at a remote location—something that is difficult to do with a graphics-intensive system interface.

Graphics-based systems present numerous advantages for use in a network “command center” environment. The latest network management systems allow operators to view the entire network on a graphics terminal. Alarm conditions produce a special color change on the graphic screen. The operator can then isolate the problem node on the screen with the click of a mouse and examine the configuration at that node. The faulty equipment component may appear in a color, such as yellow (for a minor alert), orange (for a major alert), or red (for a severe alert), that contrasts with the green of a smoothly functioning component. The operator can continue to examine different components, eventually isolating the offending card. At that point, the fault correction process can begin.

ALARM MANAGEMENT

The three main elements of alarm management are fault identification, fault isolation, and problem resolution. Effective network management systems will maximize both the capabilities of such systems and their ease of use.

In today’s T1 networks, the network has already taken care of rerouting circuits around the affected components by the time an alarm has been received at the console. Network managers must, therefore, identify and isolate problems and then track problem resolution, which is where effective network management tools come into play. The faster a problem can be isolated, the faster the network can be put back into service.

A large part of network management is trying to predict when and where future problems may occur. For this reason, an optimal system creates individual alarm records for every problem event that occurs. Alarm information is kept on-line on a hard-disk subsystem and generally is archived to tape on a monthly basis. Alarm information should be presented in an easy-to-understand format, including when and where each problem occurred, what network component was involved, when the operator responded, and what action was taken.

With a flexible configuration scheme, operators can specify the complexion of their alarm environment focusing attention on certain kinds of problems and disregarding others. Changing the severity levels of different alarms with today’s database-oriented systems is simply a matter of updating the central definitions in the alarm-reporting subsystem.

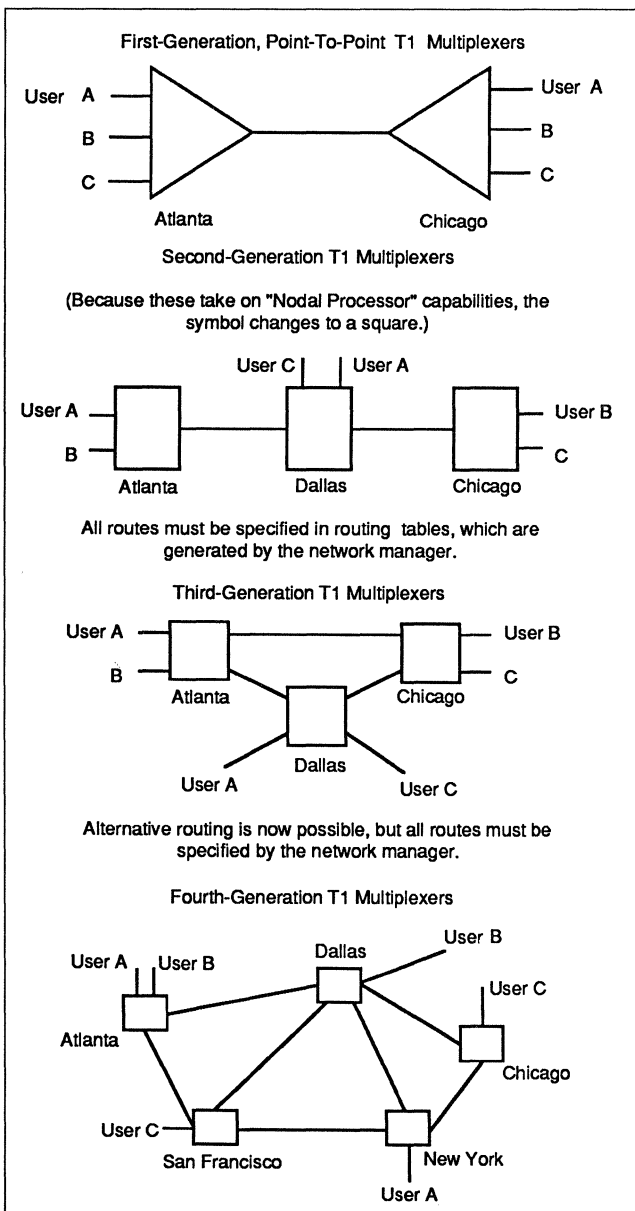


Figure 1. The four generations of T1 multiplexers.

T1 Network Management: A Strategic Perspective

CIRCUIT MANAGEMENT

Thanks to the intelligence of fourth-generation T1 multiplexers, today's T1 networks are becoming increasingly complex configurations of T1 multiplexers, nodal processors, and other "subnetworks," plus a variety of transmission facilities. In a multinode, multitrunk environment, circuit management takes on a new significance for the T1 network management system. A key development in network management—the use of relational data bases—is being used to improve circuit management. (See "The Advantages of a Relational Data Base" later in this report.)

Highly sophisticated T1 network management systems employing a relational data base allow intelligent network nodes to automatically link traffic paths between specified endpoints. These paths are configured for optimal connection based on established criteria provided by the data base's automatic information-gathering and -storage features. Circuits are rerouted and modified automatically, while the network simultaneously records and stores the information.

Despite the high degree of automation, network managers still maintain control over a certain route through the data base's two-way communication feature if specific routing criteria are necessary. This collection process enables the network manager to use the data base to analyze and solve major network configuration problems other than the normal circuit failures for which rerouting is automatically reported. In the latter scenario, circuits are reestablished to alternative sites when a single data center fails in a multiple data center network. It is the interactive dimension of the relational data base that enables the network manager to solve these problems by defining alternative configurations by specifying recovery orders. Thus, paths between nodes are reestablished with a few keystrokes, even if a catastrophic network failure has occurred. Speed is critical when such a failure happens, as reinstating the network usually requires an entire new set of off-line orders. Accomplishing this through a short set of keystrokes can mean the difference between network restoration in a matter of seconds or only after several hours.

Larger T1 networks increase the importance of *configuration management* to the network manager. Once again, it is the relational data base, with its ability to tie together a variety of network events into congruent, useful information, that enables a network manager to keep accurate, up-to-date information about the network. This includes data about each node and site in the network, including their locations, the names and telephone numbers of contact persons, the

equipment installed, and the vendor contacts—information detailed right down to the configuration of each card and port on the T1 multiplexer.

Traditional manual methods have not only presented managers with paperwork nightmares, but have created a built-in margin of error in recording data. Today's most sophisticated network management systems automate the processes of collecting, maintaining, and updating this information, so that it is available instantly whenever needed. In fact, most T1 networks today have sufficient intelligence to enable chips on cards in the T1 multiplexers to report their unique identifying characteristics to the network management system automatically. Should it become necessary for the network configuration to be changed due to the addition of new sites or a circuit outage, having immediate access to information at this level of detail greatly simplifies the network manager's job.

Another aspect of configuration management as the network grows in power and functionality is security. Because of the ease with which the entire network can be reconfigured, it is essential that the network management system implement a security system. In software-based network management systems, operators may be assigned passwords to gain access to up to eight different system levels that define functionality ranging from "inquiry only" at the lowest level to "disruptive diagnostics" at the highest level. Such a system also allows the creation of an audit trail log of all activity, including log-ons, thus making it possible to determine which operator initiated an action in the system.

THE ADVANTAGES OF A RELATIONAL DATA BASE

The key capability in any network management system is its ability to present information about the network in a clear, easy-to-understand form. Because of the large amount of information (resident or produced) in any network, two factors are especially important to the effective management of a network: the type and amount of data the system collects, and the way the information can be used or manipulated and then presented to the operator. Network management systems seek to collect and present two basic types of information:

- information about network node configuration
- information about events in the system such as alarm and performance conditions.

Continually changing conditions require network managers to be informed in real time (up to the

T1 Network Management: A Strategic Perspective

second) about problems that arise in the network. To do this most effectively, it is preferable to have real-time interactive communications between the nodes and the network management system, rather than have a polling arrangement where the network management system polls each network node periodically to check on its status. With a real-time scheme, the network management system receives information continually from the network. As the information is generated, the network management system must present it to the operator in such a fashion that it can be used to make intelligent decisions.

Because of the great deal of information received from the network, relational data base systems are becoming the preferred method of organizing and manipulating network data. Network management systems have become a strategic tool in the development and operation of networks, which themselves are becoming even more important in corporate business strategies. Network management systems can provide a great deal of extremely useful information, but only if the network operator can sort through it all. Relational data base systems easily give operators the flexibility needed to browse large amounts of data.

The volume of data in most network management applications and the complexity of relational data base applications underline the importance of processing speed. Routine searches through large data bases can easily overwhelm small, PC-based network management systems, leaving network managers waiting for results instead of acting on them.

INVENTORY MANAGEMENT

Related to configuration management is inventory management. Whereas in the recent past it may have been feasible for technicians to visit sites and retrieve inventory information or inspect or change T1 multiplexer cards, the size of many current T1 networks makes this not only impractical but cost-prohibitive. In addition, complex networks incorporate a greater amount of equipment representing a variety of vendors, which can affect reconfiguration or repairs. Managing the growing inventory of network elements thus becomes a critical responsibility. Sophisticated, software-driven systems provide a quick and efficient way to capture this information.

In addition, T1 network management systems manage more than T1 multiplexers. Large networks generally comprise a host of subnetworks, traffic from which is aggregated onto the T1 backbone. Modem

networks, multiplexer and nodal processor networks, local area and SNA networks, and others are being mingled and must work together. The question network managers are beginning to face is how to manage them all.

What is really required to manage this multivendor environment is a "single view of the network"—that is, a central point from which the network's status and activity are monitored not only from a "bird's-eye" perspective but down to a meaningful level of detail. The development of a network management architecture compatible with a variety of vendors' equipment would address this problem. By and large, each network now uses a vendor-specified proprietary software protocol. However, standards are emerging that promise some level of integration between the micro and macro worlds of networking. In the T1 area, IBM's NetView/PC is providing a common set of functions and communications protocols that allow information from several types of devices to be brought together into a cohesive network management system.

Another communications giant, AT&T, is developing a "Unified Network Management Architecture" (UNMA) and has invited selected vendors to participate in defining that architecture. Current efforts on the part of the CCITT, which is now developing a set of recommendations for constructing such an architecture based on the X.400 messaging protocol, promise to be a step in that direction. The result would eventually be a set of non-vendor-specific interface specifications for a "universal" network management protocol.

Basing the architecture on a readily recognized standard would not only increase its availability to vendors, who can write their interface specifications accordingly, but would benefit the network manager, who can then not only choose from a wider array of networking equipment, but achieve new cost efficiencies from leveraging, rather than replacing, existing equipment. AT&T's UNMA is based on such a protocol.

As promising as the multivendor solutions may be, large corporations will never be satisfied with a single network management system, and this is not the goal of an open architecture. The aim of an open architecture is to bring together pertinent data from each network so that they may be utilized in the most efficient and productive way. While such a solution is still in the future, much progress is being made in T1 network management systems today to meet this goal. □

Choosing a T1 Network Monitoring System

This report will help you to:

- Identify product features which can assist you in isolating and locating problems on T1 networks.
 - Evaluate the effectiveness of T1 monitoring systems on the market today.
 - Select a T1 network monitoring system that meets the demands of your network.
-
-

Choosing the right T1 monitoring system is not a trivial task. The result must be an operation that effectively and efficiently monitors facilities without disrupting service, even as additional equipment and services are added.

Evaluating the effectiveness of monitoring systems is even more important in today's marketplace. Competition due to deregulation has helped bring down the price of the once-expensive T1 facility and made this technology readily available to small as well as large networks. Common carriers have drastically reduced prices of T1 services, and hardware vendors continue to offer lower-cost networking equipment. The result: small and large organizations now have the capability to provide more economical, efficient, and integrated T1 service to their users.

However, this more cost-efficient technology also means dealing with a multivendor, multiservice environment. Given the complexity of today's networks, it's no wonder that most multivendor network managers experience difficulty in identifying the cause of various network problems.

This Datapro report is based on "Choosing a T1 Network Monitoring System," by Grady T. Birdsong, TelWatch, Inc. © 1988, TelWatch, Inc. Reprinted with permission.

NETWORK MONITORING FUNCTIONS

What a network manager needs, more than anything else, is a central monitoring system that can provide an overall view of the T1 network. Such a system helps the manager locate and isolate problems which, in turn, helps reduce maintenance costs and downtime, and improve the level of service to users.

At the very least, a T1 monitoring system should:

- Provide a complete listing of each type of network circuit and identify each one individually. Circuit thresholds, supervision, and cable assignments should be an integral part of the identification process, providing a comprehensive overview of all network elements.
- Monitor each segment continuously, and immediately inform the network control entity of any events which could endanger service.
- Provide event tracking by exception report for a first level of defense in maintaining a healthy network. These reports help focus on weak parts of the network, leading to a general improvement in the level of service.
- Provide diagnostic tests which quickly and easily isolate problems with supervision, levels, noise, framing, carrier losses, and bipolar variations.

Choosing a T1 Network Monitoring System

- Track problems for analyzing a sequence of events over time. A summary analysis of events, categorized by event and time, permits planned control and reliability throughout the network.
- Provide clear and concise reports for long-range planning.

MEETING THE DEMANDS OF THE NETWORK

In the multivendor, multiservice environment, a monitoring system may have a tough job in keeping watch over all the different varieties of equipment found in a typical network. The system you choose must be the focal point of the network, able to provide an instant view of any part of your network, or the network in its entirety. Specifically, today's networks demand a monitoring system that:

- Does not interfere with network functions. Out-of-service testing is not only time-consuming and expensive, but prevents a circuit from producing revenue, providing service, or being accessed by your personnel.
- Provides in-service monitoring of all circuits, but is able to identify degradations before they affect service, and should provide a hard-copy report of diagnostic and management information. This not only improves documentation of intermittent troubles, but reduces time spent in removing circuits from service in order to test them for trouble. The system

must report an event immediately, no matter what its magnitude or status, and regardless of what other activities are in progress.

- Has a distinct diagnostic capability that can analyze a circuit in real time and determine the nature of a DS1 error. A comprehensive T1 monitoring system should provide not only CRC error counts for extended superframe applications, but also information on bipolar violation rates, framing bit error rate, carrier losses, out-of-frame conditions, bit skew (slip-related), excess zeros (ones density), red and yellow alarms, and all-ones (blue, AIS alarms)—all provided with information on rate and time of occurrence.
- Facilitates a modular, cost-effective growth pattern as your network becomes larger, and is simple in format, easily understood, and fully usable with a minimum of training.

Considering all this, it's not surprising that selecting a T1 network monitoring system is not an easy task. Using the criteria discussed in this report, the network manager should be able to make intelligent choices to meet not only today's demands, but future needs as well.

Whatever the final choice, the selection must provide an in-depth account of the current level of network service, as well as information on how it might improve and develop. Ultimately, the system chosen must satisfy the one overriding concern of any network manager—cost-effective monitoring that provides maximum service and reliability. □

T1 Multiplexers in the ISDN Environment

This report will help you to:

- Prepare for the advent of ISDN.
 - Evaluate T1 multiplexer network management capabilities within the context of ISDN requirements.
 - Take advantage of capabilities provided by Fractional T1 and DACS to make a smooth transition to ISDN.
-
-

The Integrated Services Digital Network (ISDN) offers an integrated, end-to-end digital network that is capable of handling any mix of voice, data, and video. ISDN takes advantage of both circuit and packet switching (Figure 1) to provide both basic rate (2B+D) and primary rate (23B+D) services. Eventually, ISDN will include broadband services via the H0 channel (384K bps) and H11 channel (1.536M bps).

Although standards for the "H" channels are still being defined, primary rate ISDN is being implemented on the public network at an accelerated pace by AT&T, while the Regional Holding Companies are making steady progress with basic rate ISDN. Many hardware manufacturers are now working feverishly to bring ISDN out of the laboratory and into the real world to support a variety of current and emerging business applications. Consequently, the concern among network managers has shifted from wondering *if* ISDN will ever become a viable service to *when* ISDN will become available to serve their corporate locations.

This report was developed exclusively for Datapro by Nathan J. Muller. A former consultant, Mr. Muller has 18 years experience in the computer and telecommunications industries. He has written extensively on all aspects of computers and communications, and is the author of *Minimum Risk Strategy for Acquiring Communications Equipment and Services* (Artech House, 1989).

Planning Considerations

With more and more ISDN products and services becoming available, network planners must give serious attention to how ISDN can be incorporated into existing networks. Since ISDN was designed to integrate voice and data over a single, unified network, the planning process can be facilitated by recognizing that voice and data have differing requirements that are not easily reconciled, even within the framework of ISDN.

ISDN Transmission Requirements

In designing integrated networks for the 1990s and beyond, network planners must give consideration to the differing transmission requirements of voice and data. Each mode of transmission is affected by noise and delay—two critical concerns in networking.

The error rate of computer transmissions is directly proportional to the noise on the line. Errors may initiate a retransmission or trigger some form of forward error correction. Either way, time is consumed in processing, thereby reducing throughput. If the band-pass filters do not provide channels with adequate noise protection, error control will decrease and retransmissions will increase. Again, throughput suffers. Consequently, the anticipated benefits of the network will prove elusive.

T1 Multiplexers in the ISDN Environment

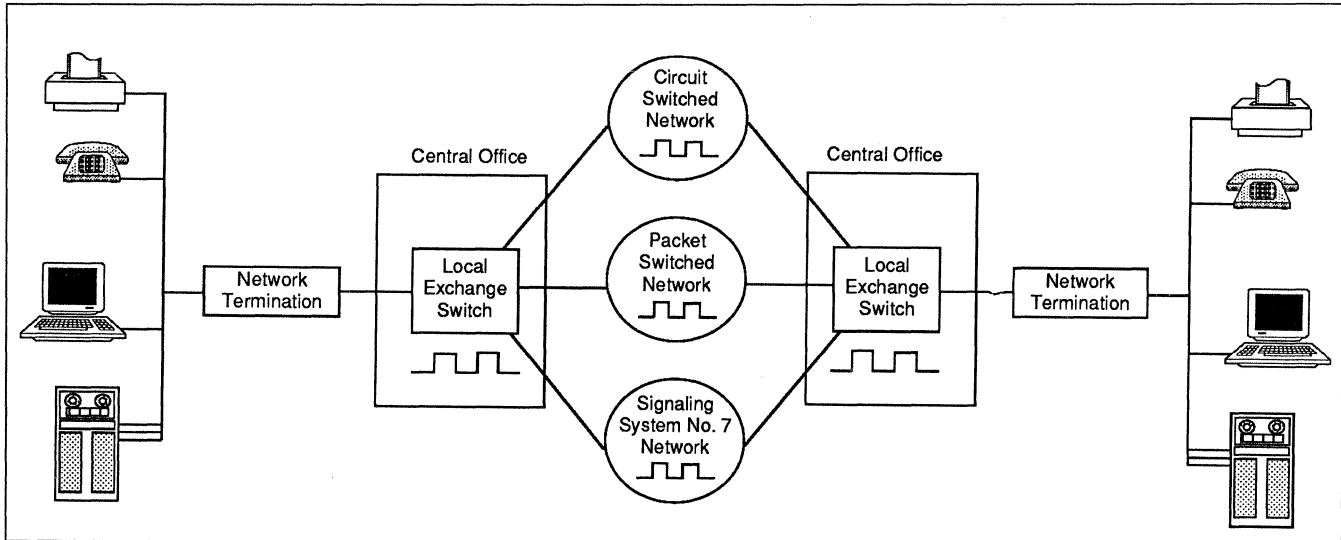


Figure 1. The evolution of ISDN.

In a voice network, the amount of delay can determine the need for echo cancelers, which are generally required whenever round trip delay exceeds 32 milliseconds. The disturbing effect (or amount of annoyance) of echoes is determined by the combination of amplitude and delay. While telephone users can tolerate a relatively high degree of delay (up to 40 milliseconds round trip delay if the echo is not too loud), this same amount of delay can cause supervisory problems for PBX tie lines, creating a condition known as *glare*. Glare occurs when both ends of a trunk are seized at the same time by different users, thereby blocking the call. Wide disparities in delay among various user locations can make glare a frequent and particularly annoying problem.

Qualitative Measurement

Delay measurement capabilities are required in multi-node, multidrop networks so that intelligent decisions can be made regarding rerouting, equipment upgrades, or contention schemes to achieve terminal-to-host response time objectives. In a multidrop network, the capability for measuring delay can be used to add precision to polling. Such precision may even permit the addition of more drop locations, while still meeting response time objectives. T1 networking multiplexers which include a software-defined transport management system can provide precision delay measurement capabilities.

The capability to monitor delay is necessary to determine the causes of changes in response time. When users order transmission facilities from interexchange carriers—or peripheral access lines (PALs) from telephone companies—they obtain the ability to commu-

nicate between two or more points. But unless the users can monitor delay, they cannot determine whether the changes in response time are the result of carrier rerouting over longer distance paths, or just too many users on the network at once.

With the ability to measure internodal delay through the multiplexer's transport management system, a user can request a better route from the carrier and monitor the result of that change. This capability constitutes a powerful tool for determining primary and alternate routes.

Data transmission also requires a transport management system that can tie back to the host. In addition to monitoring the connections between nodes, users need to know what is happening on the peripheral access lines (PALs). A transport management system can provide the capability for remotely controlled PAL restoral. This capability can reduce downtime because the PAL can be restored to service without dispatching a technician. The T1 resource manager can also remotely control the equalization of transmitters, which otherwise would require the coordinated effort of a technician at each end.

Considering the differing qualitative requirements for voice and data, the transport management system allows data to be rerouted when errors exceed predetermined thresholds. Voice need not be rerouted since it is better able to withstand impairments. There is no need to reroute all traffic, nor does a circuit have to fail before action can be taken to protect the integrity of data.

A transport management system even allows voice and data traffic to be prioritized. This capability is very im-

T1 Multiplexers in the ISDN Environment

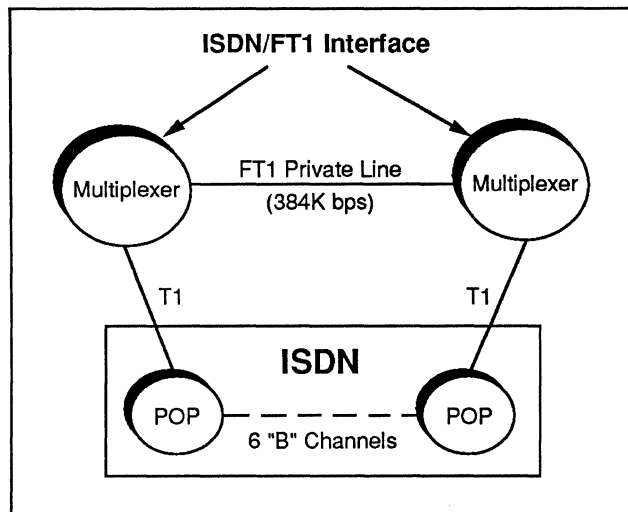


Figure 2. Dial backup to ISDN.

portant on a state network, for example, where each government agency has different requirements. State police and agencies that operate "911" emergency services have critical requirements twenty-four hours a day, seven days a week, whereas consumer agencies and motor vehicle departments use the network to conduct relatively routine administrative business only eight hours a day, five days a week. Consequently, the response time objectives of each agency are different as well as their requirements for network integrity.

On the high-capacity network, there can be several grades of service for data and others for voice. Critical data will have the highest priority in terms of response time and error thresholds and will take precedence over other classes of traffic when it comes to restoral. Since routine data can tolerate a longer response time and higher error rate, a lower priority restoral threshold is adequate. Voice is even more tolerant than data with regard to errors and delay, so restoral may not be necessary. The capability to prioritize traffic and re-route only when necessary insures maximum channel fills, which improves the efficiency of the entire network and, consequently, the cost of operation.

Traffic Activity and Density Patterns

When designing the integrated network, both traffic activity patterns and traffic density must be considered.

Because the traffic patterns for voice and data are so different, a network cannot be optimized for voice and still be valid for data. Moreover, while high-capacity transmission facilities handle a combination of voice and data, there are speeds that are unique to data. When configuring integrated networks, planners must

bear in mind that while a given transmission pipe can always handle voice, it may not handle data efficiently. Both voice and data have their own special needs.

Voice and data are also very different with regard to traffic density. In the voice world, the traffic builds up as it enters the center of the network. In a data network, the density is at one end—the host. A related consideration is the polling application in data networks, a requirement that does not exist in the voice world.

Restoral

Restoral, the concept of protecting communications pathways, comes from the computer world. Management capabilities, such as Customer Controlled Reconfiguration (CCR) available with AT&T's Digital Access and Cross-connect System (DACS), are ways to rearrange the carrier network to accommodate data or voice. The CCR may take up to a half hour or more to restore, however, depending on the complexity of the reconfiguration. Data subscribers may need to implement alternate routing more quickly. While CCR is usually acceptable as a rerouting tool for voice, it is not designed for data. Thus, while DACS/CCR can be used to remedy a long term outage, it is really not effective for instantaneous restoral.

Hybrid Networking

The concept of hybrid networking is based on a combination of private and public network elements. Hybrid networking will become an increasingly important concern in the decade ahead as ISDN becomes more generally available. The concept entails blending the services of multiple vendors over a high-capacity network. Unfortunately, it also means contending with a multiplicity of proprietary management systems from various carriers and equipment vendors.

Reclaiming control of data processing and communications resources from carriers and computer vendors is essential to achieving efficiency and cost savings. But reaping these benefits by managing one's own resources requires a high degree of independence. This degree of independence can be achieved with a transport management system that provides hooks into proprietary management systems, while offering a migration platform to the emerging Open Systems Interconnection (OSI) standard. With independence comes leverage, which can be exercised to insure the peak performance of both carriers and vendors.

T1 Multiplexers in the ISDN Environment

Other Factors

The packaging of channels onto printed circuit cards is different for voice and data. PBX manufacturers employ large scale integration (LSI) techniques to put four, eight, or twelve line circuits on a single printed circuit card. Although this packaging scheme works well for voice, it may pose problems for data. In a high-capacity nodal network, for example, a single channel may provide 32 terminals with "logical" access to a large cluster controller. In the event one channel of a four-channel card fails, three other channels must be taken out of service until a card swap can be made. If the failure occurs on the common elements of the card, a high probability exists that all its channels will be disabled simultaneously. In the application described, this failure would deny service to as many as 96 other terminals.

Discarding Tradition

Historically, carriers have paid little attention to bandwidth efficiency and transport management. Viewing themselves primarily as telephone service providers, carriers have only recently treated data transport as a major part of their business.

Despite the availability of T1 interfaces, the principal concern of PBX manufacturers has naturally been with maintaining compatibility with public network standards for voice transmission (64K bps PCM and D4 framing). As a result, PBX vendors have not emphasized bandwidth efficiency and transport management in the past—important factors when transmitting data.

Within the ISDN framework, the transport management capabilities of T1 multiplexers may offer users more control over network resources than ISDN alone can offer through the "D" channel. In fact, a transport management system complements ISDN very well, thus facilitating the implementation of hybrid networks.

Transition to ISDN

The emergence of Fractional T1 (FT1) services shows promise as a transitional offering to ISDN, especially for those who still harbor serious doubts about ISDN applications and startup costs. FT1 at 384K bps would provide a logical migration path to ISDN, since this is the bandwidth requirement for the broadband "H0" channel. Interexchange carriers already allocate 384K bps to the user. Thus, customers would require only an appropriate digital interface to derive the full functionality of the future "H0" channel when connected to the ISDN serving office. [*NOTE: Currently, FT1 is offered*

on interoffice channels only. By early 1990, the Regional Holding Companies will begin offering FT1 access to interexchange facilities through their subsidiary telephone companies. Currently, such access is provided through T1 pipes.]

DACS customers have the capability of subdividing T1 bandwidth and routing individual DS0s to their appropriate destinations. Thus, the present network architecture already provides the platform necessary for the implementation of FT1. In addition to achieving control of network load scheduling and disaster recovery with CCR, users can improve the efficiency of their networks and streamline costs by adding or deleting DS0s as needed—without incurring additional costs for altering internodal backbone facilities.

Despite the many advantages of DACS, it does not yet lend itself to realtime operations. Adding or deleting DS0s may take 15 minutes or more to accomplish, but that is still a big improvement over the current long lead times for installing new circuits that non-DACS users must tolerate.

Nevertheless, with the right networking T1 multiplexer, users would still have an economical migration platform from DS1 or FT1 and, from there, to ISDN. Primary rate (23B+D) ISDN is achieved merely by repackaging T1 bandwidth into twenty-three 64K bps bearer channels and one 64K bps signaling channel.

One-quarter FT1 equates to the 384K bps H0 unrestricted channel under ISDN. In fact, FT1 smooths the transition to ISDN in the following ways:

- DS0s may be bundled according to the bandwidth requirement of the application and routed to their destinations through instructions issued at the CCR terminal of the DACS.
- Voice and data may be integrated for transport over the same digital facility.
- Supervisory signaling required for network management may be routed through the public network to remote locations.

Despite these capabilities, FT1 differs from ISDN in some important ways. First, and most obvious, there is no separate channel available with FT1 for the transmission of signaling information. Under FT1, multiple channels including the supervisory channel would be submultiplexed into the standard DS0 (64K bps) time slots. The "D" channel under ISDN enables the user to program the network to dynamically reconfigure itself by adding or subtracting available channels and specific services on a call-by-call basis.

T1 Multiplexers in the ISDN Environment

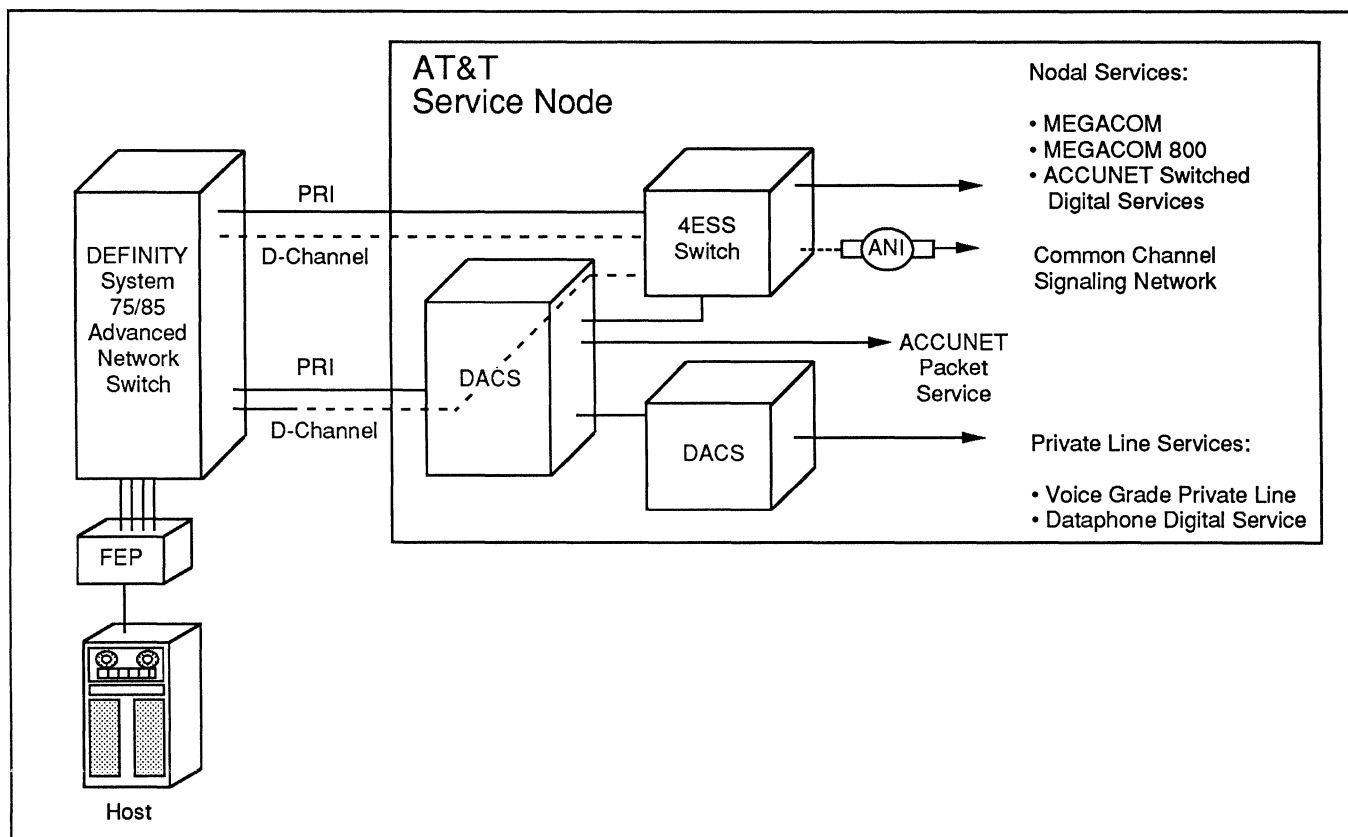


Figure 3. AT&T's ISDN PRI implementation.

Nevertheless, the DACS provides flexible bandwidth allocation via the CCR terminal. In this way, DACS/CCR may be used to rearrange and route DS0s as needed and provide access to a variety of public network services. The DACS/CCR does permit the allocation of bandwidth, but on a static demand basis. The "D" channel of ISDN, on the other hand, permits dynamic bandwidth allocation. The difference between "static" demand and "dynamic" demand is subtle, but important. Static demand entails a request for bandwidth, which may be delivered for a particular application within about 15 minutes of the request. This reflects the inherent limitation of the DACS architecture in that bandwidth requests are made via dial-up connections. Dynamic demand entails the virtually instantaneous reallocation of bandwidth on a call-by-call basis.

Despite the time limitation inherent in the "static demand" type of request, FT1 serves as a stepping stone to ISDN, offering an application-oriented bandwidth packing scheme and a reasonable amount of control.

The complementary relationship of Fractional T1 to ISDN is illustrated in the dial backup application (Figure 2). In the event the fractional link goes down,

dial backup to ISDN may be implemented in a manner analogous to that of AT&T's ACCUNET T1.5 Reserved Service. Instead of the user calling a "700" number to notify AT&T to implement the reserved circuit, the change can be implemented on a call-by-call basis through the dynamic restoral capabilities of the multiplexer under control of the transport management system. In this way, users can take advantage of ISDN on an as-needed basis and pay for only the amount of interoffice channel (IOC) bandwidth used.

IMPLEMENTING ISDN: PBX or T1 MUX

As stated previously, the network manager must carefully consider data transmission requirements when planning the integrated network. The same consideration applies to preparing for ISDN. The following paragraphs contrast the two principal implementations of primary rate ISDN—the PBX and T1 multiplexer—in light of the data transmission requirements discussed previously.

For networks that are predominantly voice-oriented, PBX implementations of PRI ISDN provide all the interconnectivity, management, and economy required for voice traffic. But for the integrated voice-data net-

T1 Multiplexers in the ISDN Environment

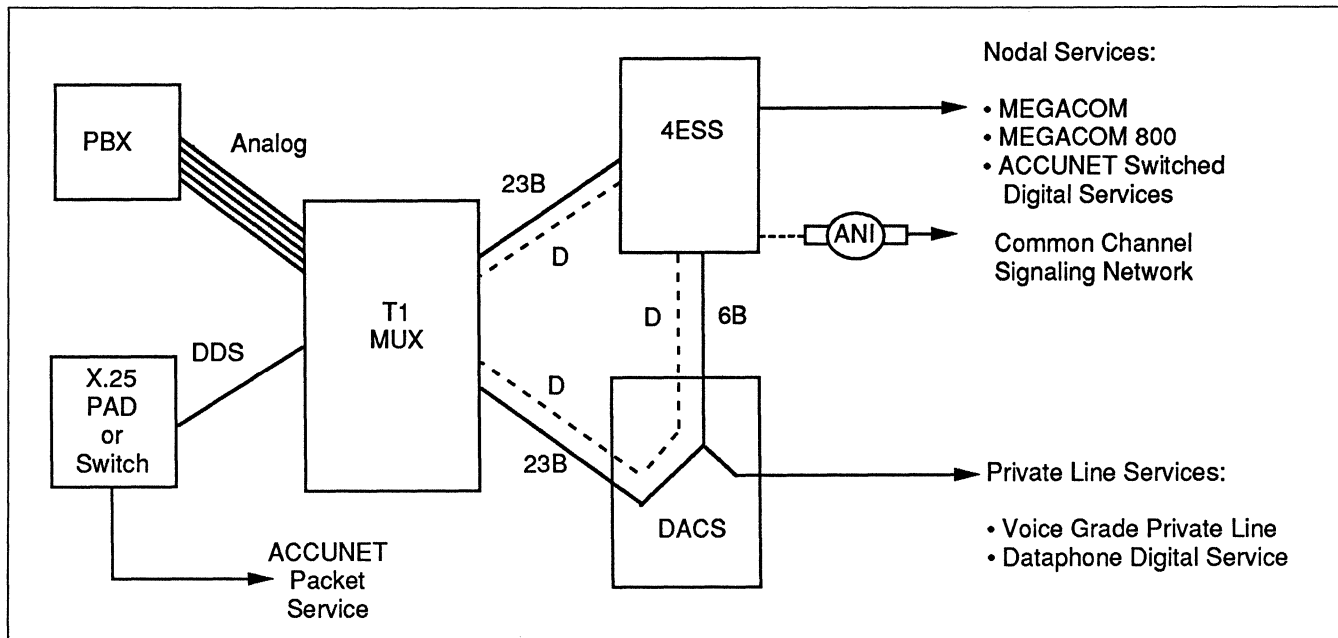


Figure 4. Multiplexer implementation of ISDN PRI.

works of the 1990s and beyond, consideration must be given to how voice and data can be managed separately over the same network, or through hybrid networks that are composed of public and private elements.

Today's digital PBXs provide direct connections to T1 facilities for trunking as well as some connectivity to SNA and X.25 networks. Many organizations continue to maintain separate networks for voice and data, however, which may indicate uneasiness with using the PBX to integrate the two. Most of this uneasiness stems from the recognition that voice and data have very different traffic patterns as well as differing transmission requirements.

ISDN via the PBX

AT&T will implement primary rate ISDN through its System 75/85 Definity PBX (Figure 3). In this scenario, digital lines (23B+D) from the customer terminate at an AT&T Service Node consisting of a 4ESS central office switch and a DACS. Access to switched services (MEGACOM, MEGACOM 800, and ACCUNET Switched Digital Services) is provided through the 4ESS. Another digital link transports voice to a DACS where the "D" channel is routed to the 4ESS via a "nailed-up" connection. The DACS routes individual voice channels to AT&T's ACCUNET Packet Service or to another DACS that provides access to private line services.

However, any PBX implementation of primary rate ISDN must take into account the following factors:

- A full 64K bps channel is required for data, whether or not that much bandwidth is actually used.
- Unlike the DACS, the 4ESS offers no DS0 "bundling" capability, which means that certain applications requiring fractional bandwidths (e.g., 384K bps) cannot be supported until the H0 channel is added to AT&T's PRI tariffs.
- There is no end-to-end network management. The "D" channel's signaling information cannot be read by the DACS as it can be by the 4ESS.
- Finally, there is no loopback capability in this arrangement and, consequently, no provision for integral diagnostics.

The PBX will likely be the principal means by which *basic rate* ISDN reaches the desktop. The T1 multiplexer, however, provides extended transport management capabilities and is likely to become the preferred method for implementing *primary rate* ISDN.

ISDN via the T1 Multiplexer

Figure 4 depicts an ISDN implementation using a T1 networking multiplexer. The multiplexer consolidates voice and data from multiple sources, such as PBXs, feeder muxes and X.25 PADs. This provides efficient utilization of available bandwidth. Multiplexing

T1 Multiplexers in the ISDN Environment

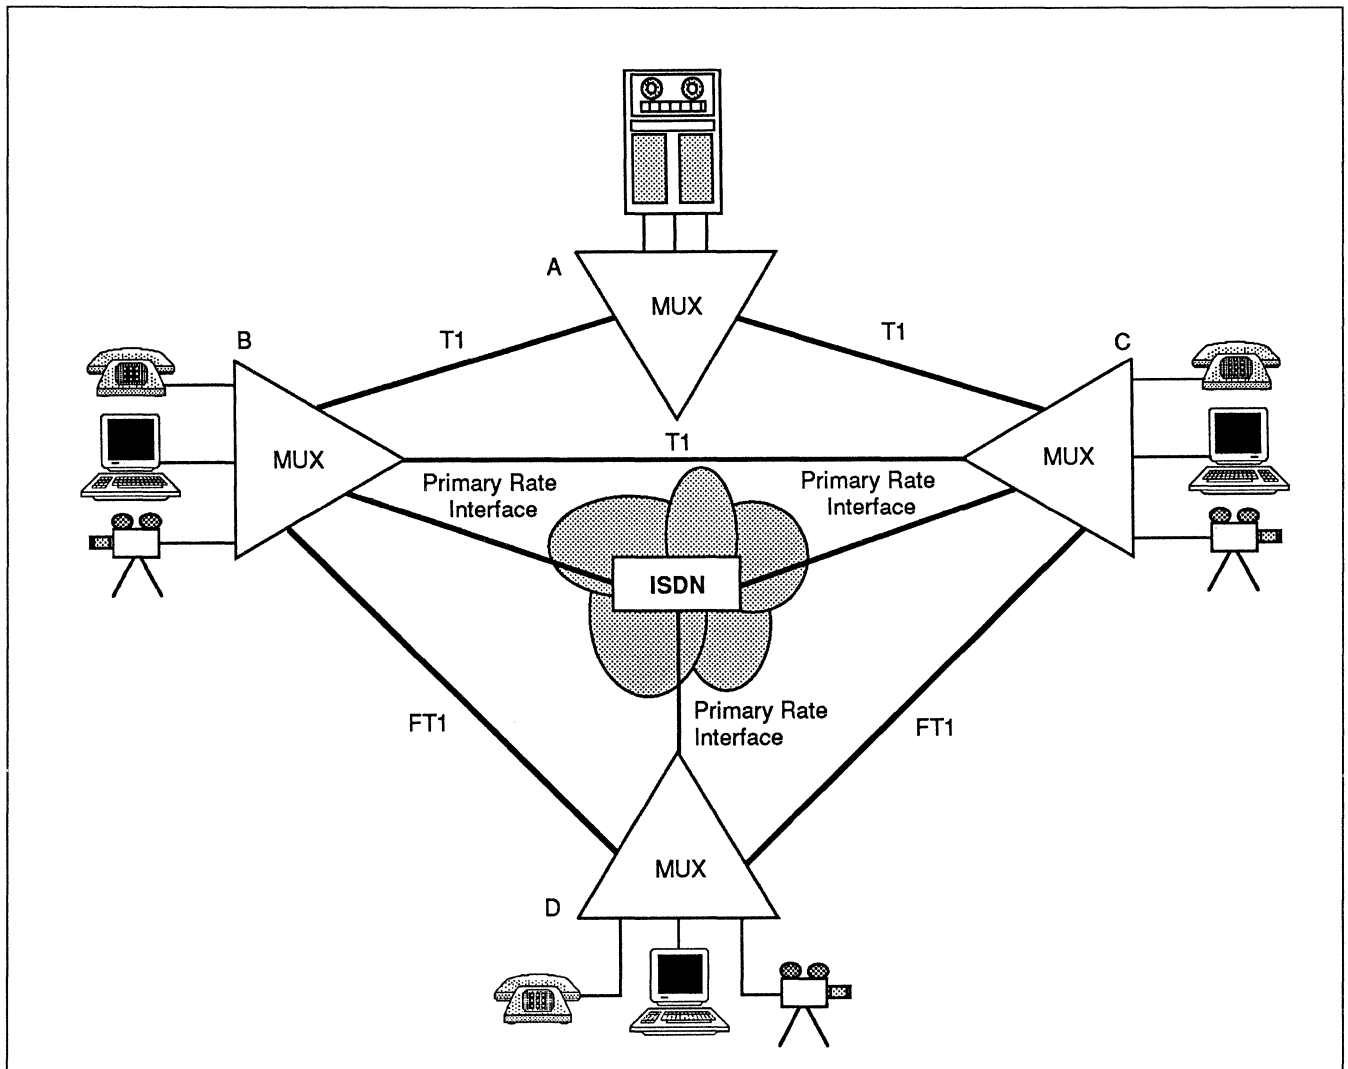


Figure 5. Hybrid networking with FT1/ISDN.

within DS0's or "B" channels squeezes multiple voice or data channels into a single 64K slot, adding even greater efficiencies to bandwidth utilization. Applications requiring fractional bandwidth can be passed through the 4ESS once appropriate tariffs are implemented.

Since the 4ESS is capable of accepting fractional bandwidth at 384K bps (six contiguous DS0s), it can be used for disaster recovery. And since the multiplexer can embed network management information into each of the DS0s or DS0 "bundles," end-to-end management and control is preserved even within the ISDN framework. This means that integral diagnostics are possible, since loopback testing can be performed.

Implementing primary rate ISDN through the T1 networking multiplexer permits the integration of voice and data with efficient utilization of available bandwidth—without degrading the traffic handling ca-

pabilities of the PBX. Bandwidth can be fractionalized according to application, instead of limited to increments of 64K bps under ISDN. Moreover, the user can realize the benefits of integral diagnostics and preserve the integrity of network management end to end.

Although several multiplexer vendors have announced "ISDN-ready" or "ISDN-like" capabilities, they offer only passive conveyance of PRI. True ISDN functionality consists of the intelligent linkage between the transport facility (T1) and premises equipment. This means that "B" channel routing must take place under "D" channel control between directly connected T1 multiplexers and ISDN PRI switches (i.e., 4ESS or 5ESS).

T1 Multiplexers in the ISDN Environment

HYBRID NETWORKING

Despite the growing appeal of ISDN for integrating voice, data and video, it will not cause corporations to suddenly abandon their investments in private networks. ISDN is a service offered through the public switched network; as such, users will have to give up the level of control they now have with private networks.

Corporations will more likely maintain private networks when it is advantageous to do so, and use ISDN when it is advantageous to do so, just as they now use other public network elements such as DACS/CCR, M24/M44, or SDN. Blending public and private network elements in this way is referred to as "hybrid networking."

The next issue, then, is public versus private networking. Each approach, of course, has its advantages. But the choice of one over the other is not so clear cut. Users need the efficiency and control of private networking along with the connectivity and access to services that is associated with public networking. Therefore, it would be advantageous for most users to position themselves to implement the appropriate network architecture (public/private) to realize the best return on investment at any given time—even when ISDN becomes universally available.

But, mixing private and public network elements into a hybrid arrangement is not enough. The ability to implement network management and control, regardless of traffic type, is equally important.

A T1 multiplexer capable of embedding supervisory signaling into the "B" channels can provide interoperability with both proprietary and DS1/DS0-framed products. This capability can assist users in building hybrid networks composed of private and public elements, and improve ISDN's efficiency and economy.

Enhanced ISDN

Improving the efficiency and economy of ISDN would entail the multiplexer treating the "B" channels of ISDN as "subaggregates." When "B" channels are routed by the "D" channel through a public network element such as the 4ESS, multiple voice or data chan-

nels and supervisory information can be embedded within each 64K "B" channel. Since the 4ESS does not route the "D" channel, embedding supervisory information within the "B" channel provides users with the means to extend private network management and control through the public ISDN network. That way, users do not have to sacrifice control for the sake of cost savings.

Another way that such a multiplexer can enhance ISDN is by enabling the "B" channels to carry ADPCM compressed voice at 32K bps, 24K bps, and 16K bps. Software-controlled speed selection gives users added configuration flexibility that does not come with ISDN alone, since ISDN uses only 64K PCM voice. Although the current version of the ISDN standards specify the entire "B" channel as the fundamental unit of circuit switching, logical subchannels may be carried on a single circuit between the same pair of subscribers.

As AT&T expands PRI tariffs to embrace 384K bps "H0" channels and 1.536M bps "H11" channels, the multiplexer will provide users with the means to implement "wideband" dynamic disaster recovery. Well-designed multiplexers will offer this "wideband" function as a software-expandable option. In this case, the initial ISDN hardware and software must be designed for future "H0" and "H11" tariffs without requiring a "forklift" conversion (or upgrade).

For those who cannot yet justify ISDN on the basis of cost or applications, the multiplexer's DACS-compatibility and support of FT1 would permit flexible bandwidth allocation, integrate voice and data, and extend network management and control to remote locations. The PRI can be added later when needs justify.

With FT1, users can actually have more flexibility in sizing bandwidth to fit specific applications, since primary rate ISDN—as currently tariffed without the "H0" channel—is limited to providing 64K bps "B" channels, which cannot be bundled to meet data application needs. In supporting both ISDN and FT1, a T1 multiplexer—with the requisite functionality embedded in its transport management system—further extends the options available to users when building hybrid networks (Figure 5) □.

Local Area Network Management Issues

This report will help you to:

- Identify the facilities required to manage local area networks.
 - Apply a framework which relates LAN management to the OSI architecture.
 - Understand how OSI management protocols, services, and layers interact.
-
-

The ISO Reference Model for Open Systems Interconnection (OSI)¹ provides a framework for the description and standardization of communication activities. It roughly distinguishes between “normal communication” and “management” activities. Normal communication includes all functions which concern themselves with the actual transfer of data, while management has become the collective term for all activities concerned with planning, organizing, supervising and controlling the communication system.

This report specifically addresses the management of local area networks in closed distributed computer systems. It presents the application and administrative environment and the network characteristics from which the LAN management requirements and functions are derived. Several methods of classifying management activities are discussed. In addition, the report presents a management framework which relates management to the OSI architecture and a Resource Management Model that defines the interactions between the communication and the management entities as well as between the management entities themselves.

Of the concepts presented in this report, many are taken from a report⁵ on a collaborative project funded under the COST-11 bis programme by the Commission of the European Community. Others review standardization activities in ISO (TC97/SC16), IEEE 802 and ECMA (TC23/TC24).

LAN MANAGEMENT CONCEPTS

This section presents the application environment and the network characteristics from which the LAN management requirements and functions are derived.

LAN Management Environment

Application and Administrative Environment. The report specifically addresses the management of closed distributed computer systems. These systems are usually dedicated to one particular application and are designed to meet its special requirements. They are often restricted to a closed user group and not open to users and applications from outside. Communication systems, application and operating systems are closely integrated and are under the control of one single organization. Typical application areas are process control, data acquisition and factory automation.

The networks used in such applications do not quite correspond to the ISO view of “Open Systems Interconnection (OSI).” In the ISO view, an open system

This Datapro report is based on “Local Area Network Management Issues,” by Voler Tschammer & Horst Klessman. Excerpted from *EFOC/LAN '85 Proceedings*, published by Information Gatekeepers, Inc. 214 Harvard Ave. Boston, MA. 02134 (617) 232-3111. © 1989 by McGraw-Hill, Inc. Used by permission.

Local Area Network Management Issues

is an autonomous entity and the OSI standards primarily cover the aspects of heterogeneity and geographical separation.

In closed systems the interconnected computers and devices are not completely autonomous but cooperate closely to achieve a common task. This results in intensive communication activities which are subjected to particular requirements on performance, reliability and flexibility. These are usually not common to OSI applications.

Network Concept. Local area networks usually have a simple network topology such as a ring, a star or a bus with all stations sharing a common transmission medium. This allows communication by information broadcast and enables each station to continuously monitor all activities on the communication channel. Broadcast communication allows direct information exchange between any two stations providing a basis where upon one-to-one, one-to-many (multicast, broadcast) and any-to-one (client/server type) information flow patterns can easily be implemented. In such networks special promiscuous stations may be installed, which receive all messages from the LAN. These stations passively monitor the network traffic and collect statistics.

LAN Management Requirements

Network management in distributed computer systems concerns activities which allow the system to satisfy the requirements on performance, reliability, availability and flexibility imposed by the application. Management entities perform all functions necessary for planning, organizing, supervising and controlling the components of the system. LAN management observes the operation of the network, makes decisions based on these observations, and invokes control actions which support the correct network behavior and handle errors. Thus, LAN management is primarily provision, maintenance and optimization of communication services and network performance.

User administration seems to be a secondary management function in closed systems. Access control, security, protection, accounting and billing are very important when the network is under the control of an independent organization which charges users for the utilization of network services. This is usually the case with wide area networks. In closed systems the users of communication services are normally under the control of the same organization as the network itself. Users tend to be devices or programs rather than humans. Human users are operators or engineers who interact with the system via well-defined

application programs. Thus, user administration is primarily a function of the application and not of the communication system.

Management Functions

This section identifies the facilities required for the management of local area networks. They are derived from the management environment and the requirements presented in the previous sections. LAN management must support the normal communication activities; handle errors and faults; install and modify components; and monitor, analyze and optimize system performance. These activities may be classified into the following categories:⁵

Operational Management. Operational management functions support the normal communication activities and are involved in servicing requests from higher layers. They correspond to the control and supervisory functions found in most communication protocols and are usually specified and implemented with normal communication functions and protocol entities. They are responsible for the maintenance of status, mapping and routing information as well as for allocation and management of buffers, virtual circuits, transmission media and other communication resources.

Maintenance. Faults and serious errors must be reported to the management which is responsible for error logging and recovery. Management entities must involve diagnostic tests which locate a fault to the level of a "smallest replaceable unit" which must be replaced or repaired. Management entities must also test new or repaired components to determine whether they meet their functional specifications and to ensure that they do not disrupt the operation of the network when they are put "in service."

Configuration Management. Configuration management controls the installation and modification of hardware and software components. This is necessary for reconfiguration after failures and in order to allow flexibility. Basic reconfigurations are: bypassing bottlenecks and faults, activating backup resources, updating routing tables, reassigning functions, and reallocating resources such as buffers, virtual circuits, etc. These activities are driven by events which signal a degradation of service quality or a component failure caused by internal errors or external influences. Flexibility allows new functionality and new technology. Components are replaced to meet changes in networking techniques and application requirements.

The installation of new components may require specific down-line load and bootstrap facilities. The

Local Area Network Management Issues

proper allocation of addresses and names to components is also a configuration activity. Likewise, the management must log configuration changes and maintain up-to-date status information on the hardware and software components, including version number, current state, connectivity, etc.

Performance Management and Optimization. Many LAN applications need to know what performance to expect from the network. Usually this is difficult to predict and therefore must be measured under operational or test conditions. Traffic statistics and error counts must be collected by the network management and subsequently processed and analyzed to detect bottlenecks and negative trends such as decreasing throughput or increasing transfer delays.⁷ A global view of the network performance is necessary to detect design faults and to allow for optimization by tuning parameters, dimensioning resources or replacing components.⁶

User Administration. Management activities concerned with user administration are very important when the network is under the control of an independent organization which charges users for the utilization of network services. This is usually the case with wide area networks. In local area networks, particularly in closed systems, the users are often devices or programs under the control of the same organization as the network itself. Thus, user access control, security and accounting, rather than being within the responsibility of the communication management, are part of the application.

LAN MANAGEMENT ARCHITECTURE

This section presents a framework for structuring and classifying LAN management functions and interactions. The OSI Reference Model is reviewed and differences between the OSI management concept and LAN management structures are elaborated.

OSI Management Framework

ISO has recognized the need to include management aspects in the OSI Reference Model. However, according to OSI principles, it has restricted its concern to those activities which imply actual exchanges of information between open systems. In this context it identifies the following categories of management activities:

Application Management. Application management relates to the management of OSI application processes. Typical activities are initiation, maintenance and termination of application processes, manage-

ment of resources, resource interference and deadlock prevention, integrity and security control. The application management protocols reside within the Application Layer.

Systems Management. Systems management relates to the management of OSI functions and resources across all layers of the OSI architecture. Typical activities are activation, maintenance, and termination of OSI resources, connection management, performance monitoring and analysis, reconfigurations and error control. The systems management protocols reside within the Application Layers.

Layer Management. Layer management has two aspects. One concerns layer activities. It is implemented by the layer protocol to which it applies. The other aspect is a subset of systems management. Its protocols reside within the Application Layer and are handled by the same type of entities which handle systems management protocols.

Positioning Principles. Both centralization and decentralization of management functions are allowed. This calls for a structure in which each open system is allowed to include any subset of systems management and layer management functions. Associations between management entities can be established at any time, in particular at the time when a system which has been operating in isolation from other systems becomes part of the OSI environment.

LAN Management Framework

ISO and several other standardization bodies (ECMA, IEEE 802) work on the basis of the OSI management concept. The aim of their activities is to provide a framework for the interconnection of systems management application processes.

ISO working papers² identify two aspects of systems management. One aspect is to coordinate the activities of the seven layers in a given open system; the other aspect is to coordinate the activities of the various interconnected systems. The first aspect must be covered in every open system, the second aspect is optional because a manager may or may not communicate with other managers. The first aspect is further called "intrasystem responsibilities." It has only local significance and is therefore declared to be outside the scope of OSI. Only the second aspect, called "intersystem responsibilities," is inside the scope of open systems management, because only these activities contribute to the coordinated management of interconnected open systems. This can be summarized as follows:

Local Area Network Management Issues

1. The cross-layer aspect of management has only local significance and thus is outside the scope of the OSI management architecture.
2. The network-wide aspect of management is within the intersystem responsibilities of particular interconnected application processes, which are called systems management application processes.

This view does not completely satisfy the management requirements of local area networks in closed distributed systems. The reasons can be exemplified by the differences between Network Operating Systems and Distributed Operating Systems.⁸

In the Network Operating System (NOS) approach, each interconnected computer continues to run its local operating system which is independent of the network. The NOS is built on top of existing operating systems and implemented as a collection of application processes. The NOS is responsible for communication and resource sharing and attempts to hide the differences between the underlying systems. In many respects this is identical to the OSI management architecture.

The Distributed Operating System (DOS) approach refers to an integrated network with one homogeneous operating system for all distributed computers. Usually the DOS is designed for one particular type of application and with its requirements in mind. DOS functions are replicated in each network node and manage the operations and resources of the distributed system in a uniform and global fashion. The same approach must be taken in order to satisfy the LAN management requirements in closed distributed systems. The differences to the OSI management architecture can be summarized as follows:

- The cross-layer management aspects are not restricted to local significance only. They must be included in the LAN management architecture.
- The intersystem responsibilities are not restricted to interconnected application processes. Every management function on each layer may have network-wide responsibilities and may communicate directly with peer-management entities.

This approach is represented by the LAN management model described in the following section.

A LAN Management Model

The Basic OSI Reference Model. The LAN management model was developed on the assumption that the OSI Reference Model already represents both the

communication and the management functions. Thus, by applying the same abstractions, that were used for the development of the OSI model, it should be possible to further refine the OSI model to represent the LAN management structure and the relationship between management and communication entities. The principles used in the development of the OSI model were geographical separation and hierarchical order. These principles can be regarded as a "horizontal" and a "vertical" refinement to the most general situation of application processes interacting via a communication system.

Geographical separation represents the fact that users may be located within different, geographically separated stations, which are connected via a transmission medium. The interaction between the stations is represented by the peer-to-peer communication which is governed by protocols. Layering represents the hierarchical order of communication functions. Each layer performs a specific set of functions which add on or enhance the functions performed by the lower layers. Each layer provides a specific set of services which add value to the services provided by the lower layers in such a way that the highest layer is offered the set of services needed to run distributed applications. Consequently the interlayer relations are represented by abstractions termed "services."

The combination of geographical separation and layering leads to the well-known OSI Reference Model.

Management Refinement. Management is represented in the model through a further step of refinement. It distinguishes the pure communication functions from the communication management. The interactions between both types of entities must be further defined. One method of defining these interactions is introduced by the Resource Management Model described in a later section. This model considers the management a process control system which monitors and controls the communication entities and resources.

The refinements of layering and geographical separation as found in the OSI model can also be applied to management. This leads to the layered, distributed management structure which is illustrated in Figure 1.

It includes management protocols representing the geographical distribution and significance of management activities, and it includes management services representing the hierarchical order of management functions.

This management model will serve as a framework for further refinements of the LAN management ar-

Local Area Network Management Issues




Figure 1. A LAN management model.

chitecture. The first steps of refinement should consider a detailed definition of the interactions between management and communication entities (monitoring, control) as well as between the management entities themselves (management services and protocols). The hierarchy of management functions, i.e., the proper definition of management layers, is another open question.

LAN Management Structure

Management Layers. A first and obvious approach to the selection of specific management layers may be based on the seven layers of the OSI Reference Model. Figure 1 illustrates this layering of management functions. The structure represented by this figure also comprises a nonlayered management component. This reflects those management activities which cross layers. The cross-layer component is included in the LAN management architecture and is not restricted to systems management application entities. This allows a management architecture which implements the DOS approach for closed systems rather than the NOS approach usually implemented in the OSI environment.

Another management layering approach, described below, is more applicable in regards to the particular conditions of LANs in closed distributed applications:

- It is realistic to make allowance for the fact that there are different types of LANs supplied by manufacturers or organizations. Each LAN is characterized by a certain transmission medium (coaxial cable, fiber optics), topology (star, bus, ring), access scheme (CSMA/CD, Token Passing) and low-level protocol (HDLC, datagram). Management facilities will always reflect these characteristics. This leads to the identification of a "Subnet Management" layer as the lowest layer of a LAN management architecture. Further refinements may separate this layer into a Media Access and a Logical Link management sublayer.

- It is necessary to manage the communication activities across the boundaries of interconnected LANs. Management is responsible for the provision of service qualities and network performance despite different networking techniques (access schemes, protocols, etc.) and different administrative environments. This leads to the identification of an "Internet Management" Layer as the second LAN management layer.

- The management of the communication between end users is the last function which ends the management of a complete Transport System. Different host characteristics must be handled to satisfy application requirements for determinism, reliability, availability, etc. This leads to the identification of a "Transport Management" Layer as the third layer of a LAN management architecture.

- Directly above the Transport Management, the Application Management concerns itself with application-oriented communication activities. This reflects in particular those LAN architectures where the entire network is operated by one single organization, the representation of data is the same throughout the network, and application entities have direct access to Transport Services for performance reasons. For use in other application areas and nonhomogeneous networks the Application Management Layer must be subject to further refinements in order to represent activities concerned with nonhomogeneous data representation, user access control, etc.

Management Classes. In addition to geographical separation and layering, a further refinement classifies management in terms of real-time constraints.⁵ This makes allowance for the promptness with which management interacts with communication entities. Three classes may be identified:

1. Short Term Management (STM) functions are embedded in the communication entities and operate synchronously to the execution of communication

Local Area Network Management Issues

functions. They support the normal operations and are involved in servicing requests from the higher layer. STM functions correspond to the control and supervisory functions found in most communication protocols. In the management model they are represented by the communication layer entities.

2. Medium Term Management (MTM) functions are not embedded in processing communication requests. They run asynchronously with respect to communication entities and are driven by events which signal a degradation of service quality or a failure caused by internal errors or external influences. MTM functions are also constrained by the promptness of execution because their objective is to keep the communication system in service. In the model most of the MTM functions are represented by the layer management.

3. Long Term Management (LTM) functions are primarily concerned with planning, evaluation and optimization of services. LTM functions realize long term strategies and are not bound by real-time constraints. LTM functions are generally cross-layer and as such can be associated with the systems management component of the model.

Management Interactions

This section describes the interactions between management and communication entities as well as between the different management entities themselves.

Management Protocols. Management protocols represent the peer-to-peer communication between management entities. There are several ways of exchanging management information between peer-management entities situated at different locations.

In the OSI environment, management information exchange across system boundaries is primarily between systems management entities. These entities reside within the Application Layer and thus are able to exploit the full set of facilities and services available from the seven layers of the OSI architecture. Management information may be transferred as transparent user data and/or as special protocol data embedded in the normal communication protocols. The communication services will allow for heterogeneity, security, integrity, etc.

In the closed system environment, dedicated management protocols are required within all layers. One reason, which is also discussed in the OSI environment, is to maintain the management information flow despite failures or other adverse conditions, which block or intercept the normal communication path. This requires additional or alternative paths for

management data which allows management protocols to be carried directly by any layer of the OSI architecture.

The Unidata Service³ proposed by IEEE 802 may be very useful in supporting these management information paths and protocols. It is an independent, self-contained data transfer operation, which is independent in the sense that there is no relationship to any other data transfers and self-contained in the sense that all protocol information required to deliver the transfer is presented to the service provider, together with the data, in a single service access. In the LAN environment, this will probably be a connectionless service, capable of supporting multicast facilities. It is obvious that management dialogs conducted via lower layers will be less reliable than those conducted via higher layers. However, if parts of the communication system are inoperable or are approaching inoperability, the risks of low-quality services will be accepted provided that they maintain any management information path at all.

Another reason for the necessity of lower layer management protocols is to support Short Term and Medium Term management functions and provide management services to higher layers. A typical example is the Media Access Protocol in LANS. In LAN topologies where all stations share a common transmission medium, a common consensus is needed about how to access the communication channel, how to avoid or to resolve collisions, how to recover from protocol errors, how to add and remove stations, etc. This is achieved by exchanging dedicated management information, e.g., token passing, and is completely within the responsibility of the subset manager. No higher layer functions are factually involved in the media access scheme except for Long Term activities such as status monitoring and statistics in order to detect design faults and to optimize the performance.

Management Services. Management services are abstractions which represent the hierarchical order of management functions. In the management model each service offered by a layer has a communication and a management component. Management services can be defined according to the same principles used to define the communication services:

“Each layer wraps the lower layers and isolates them from the higher layers.” For example, each of several different media access control schemes, which may exist in a network of interconnected LANs, will be within the responsibility of a separate subnet manager. If LANs are replaced, higher layer management functions, e.g., end-to-end management, need not be affected.

Local Area Network Management Issues

“Each layer performs a specific set of functions, which add on or enhance the functions performed by lower layers; each layer adds value to the services provided by the set of lower layers.” An internet manager for example relies on the combined services provided by the subnet managers of the interconnected LANs. The transport manager provides end-to-end communication management, relying on the services provided by the internet manager and the managers of different resources (interfaces, communication software, buffers, processors, etc.) within the interconnected stations.

There are also service type interactions between the layered and the nonlayered part of the management structure. Long Term management is generally cross layer, implementing long term strategies and policies, which are performed by invoking Medium Term management operations. For example, a cross-layer management function such as accounting, may require the correlation of statistics accumulated within several layers. These statistics in return will rely on counters and status information provided and updated by layer-specific management functions.

Layer Management Interactions. The interactions between layer-specific management functions and communication entities must be further elaborated. One method of defining these interactions is introduced by the following Resource Management Model.⁵ This model considers the LAN management a process control system which monitors and controls the communications entities and resources. Thus, the layer management interactions and relations are represented by abstractions termed “monitoring and control.”

The model relates the state and the availability of communication services to the state of resources, which are used to provide the service. The resources are within the responsibilities of the communication management. Resources of different levels of detail can be identified depending on the view of system structure. Typical views consider the logical or functional structure, the implementation or a time-dependent management structure. Stepwise refinement allows each view to identify different levels of abstractions. Thus, a resource may be viewed as a simple component at one level of abstraction although it consists of several nested resources which are physically distributed within the network.

Each resource has a state indicated by a state variable. Resource states (see Figure 2) identify whether a resource is “in-service,” or whether it is “out-of-service.” Refinements to these states may reflect different qualities of service, various transitory errors and various out-of-service states, which are under-

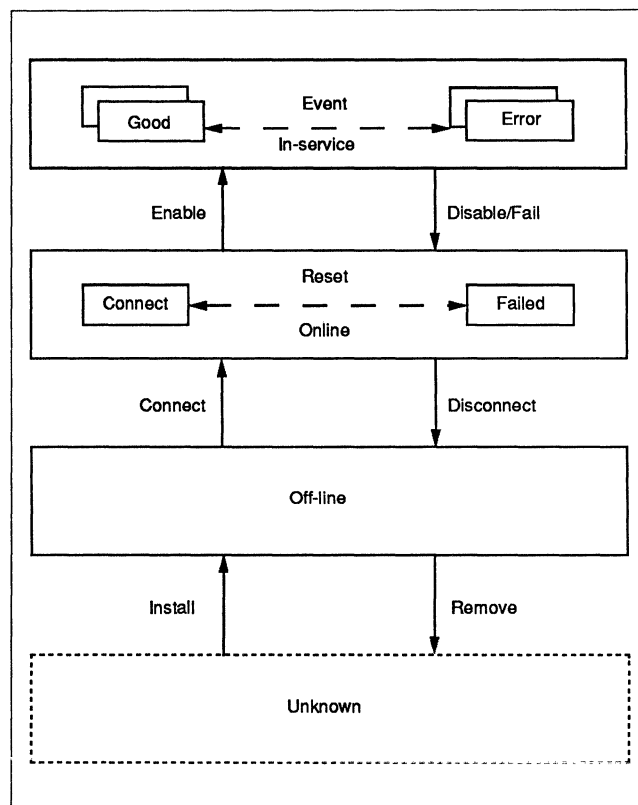


Figure 2. Resource state transition diagram.

gone when a resource is replaced or switched into a maintenance state. A resource in the in-service state works and contributes to the service provided. An on-line/out-of-service resource is physically and logically connected into the system but does not provide a service, i.e., it is in a failed state or a maintenance state. A resource in the off-line/out-of-service state has been installed but has not been connected into the system, i.e., it is not able to provide a service. This may be an intermediate state of the installation process or an off-line maintenance state. State transitions result from events, i.e., internal errors or external events, or from management actions. Management operations are classified into monitoring and control.

Monitoring allows the management to observe resource states, setpoints and parameters. Operations of the type “request” enable the management to read the values of these variables, while operations of the type “indication” allow the resource to asynchronously indicate a change of state. The generic monitor operations read-state-variable and read-set-point are able to satisfy the requirements of various management functions, such as status monitoring, events logging, error reporting, diagnostics, fault location, performance monitoring, etc.

Local Area Network Management Issues

Control operations allow the management to change resource states directly, or to control or tune the resource by changing the value of setpoints. Connect/disconnect, enable/disable and reset are used to change the state of a resource between off-line, on-line and in-service states. Install and remove operations introduce a resource to the management system or delete its name from its manager respectively.

These generic control operations are able to implement various management functions, such as initialization and close down, software load and distribution, error handling and various reconfigurations.

CONCLUSIONS

The report describes a management environment for LANs in closed applications. It outlines differences to the OSI management framework and pleads for a dedicated management architecture. The LAN management approach favors an integrated homogeneous management structure which monitors and controls communication activities and resources in a global and uniform fashion.

A LAN management system resembles a distributed control system and must be designed as an integral part of the communication system with the application requirements and the network characteristics in mind. Intersystem responsibilities are not restricted to interconnected system management application

processes, as found in the OSI environment, but are a characteristic attribute of all LAN management activities.

Management structure and interactions have been introduced according to various models and methods. Further elaboration is necessary to investigate management structures, which operate under real-time constraints and in the presence of faults and which are able to actively control and optimize the network performance.

REFERENCES

- ¹ISO 7498, "Information Processing Systems—Open Systems Interconnection—Basic Reference Model," ISO Central Secretariat, CH-1211 Geneva 20.
- ²ISO/TC97/SC21/WG16-4, "Management Information Service—Part I—General Description," doc. Paris-46, Secretariat USA (ANSI), Paris, February 1985.
- ³IEEE Project 802, "Local and Metropolitan Area Network Standards," Draft IEEE Standard 802.1: Section 5, Systems Management, Rev. E, August 1984.
- ⁴ECMA/TC23/84/154, ECMA/TC24/84/207, "OSI Management Architecture," Introductory Note to Draft TR, December 84.
- ⁵COST 11 bis Local Area Network Project, "Management of Local Area Networks," Part 2 of Final Report, edited by M. Sloman, Department of Computing, Imperial College, 180 Queensgate, London, UK SW7 2BZ.
- ⁶M.D. Abrams et al., "The NBS Network Measurement System," IEEE Transactions on Communications, vol. 20, no. 10, October 1977, pp. 1189-1198.
- ⁷R. Brice, W. Alexander, "A Network Performance Analyst's Workbench," ACM Comp. Network Performance Symposium, Performance Evaluation Review, vol. 11, no. 1, April 1982, pp. 138-146.
- ⁸A.S. Tanenbaum, "Computer Networks," Englewood Cliffs, NJ: Prentice Hall, 1981. □

Token Bus/Ring LAN Management Concepts and Architecture

This report will help you to:

- Understand token bus/ring LAN management methodology.
- Evaluate centralized vs distributed management.
- Confront implementation issues in performance, fault, and configuration management.

Management of network resources is currently one of the most critical areas of concern in local area networks (LANs). Greater reliance in the network services places tremendous importance on the network resources to operate correctly, continuously, and according to design goals. Any problem associated with the normal network operating condition must be identified, isolated, and corrected. A problem, in this sense, is an erroneous condition which results in a loss of system availability or a service deprivation to end users. As the LANs grow to accommodate more and more users and resources, and as they evolve into more complex interconnections of multiple networks, the probability of occurrence of such problems proportionally increases.

Network management is intended to control this complexity by providing request/response tools for communicating management entities at each network station. A manager station acts as a report collection and action distribution node which the other stations report to. Management request/response services between the manager and the agent stations are being described in management standards. These standards do not indicate, however, how the management of a certain function is accomplished or implemented.

This Datapro report is based on "Token Bus/Ring Local Area Network Management Concepts and Architecture," by Tuncay Saydam and Adarshpas S. Sethi. © 1987 IEEE. Reprinted, with permission, from *IEEE Infocom '87 The Conference on Computer Communications Proceedings*, San Francisco, CA, March 31-April 2, 1987, pp. 988-993.

This report will emphasize the LAN management architecture, management modes and implementation issues.

PREVIOUS WORK ON LAN MANAGEMENT

The area of network management research is very young, only a few years old. One may view the LAN standardization work by IEEE 802 committees as one of the first research efforts. Other organizations (ECAM, ANSI, ISO, GM(MAP)) have also proposed LAN management architectures as extensions to current standards. However, the development, specification and implementation of actual LAN management requirements fall behind these standardization efforts. It seems that the interest in developing management systems exceeds the user specification of their proposed uses.

IEEE standards 802.4¹, 802.5², 802.1 Part B³, and 802.3⁴ on systems management in LAN environments contain a rather complete description of a management architecture and protocol tools necessary to manage LLC, MAC and physical layers. Thompson⁵ is a clear summary of management architecture, services and protocols. Carlos and Winkler⁶ define three levels of management hierarchy before token-ring LAN management. The network manager acts as an overall controller while management "servers"—three different types—act as data collection and distribution points which the stations report to. Useful for multiring networks, it appears to waste

Token Bus/Ring LAN Management Concepts and Architecture

bandwidth, increases management traffic for single LANs as well as introduces added complexity.

There is a growing interest in research on LAN management and we expect to see an increasing number of publications in the near future.

MANAGEMENT STANDARDS

There have been ever increasing efforts to develop network standards. IEEE and ISO committees have been addressing the management issues and producing the key standard documents 1, 2, 3, 4, 5, 6, 7. These documents provide a set of management tools for system management operations. They do not prohibit the use of management mechanisms other than those provided in the protocol and do not constrain implementation except as necessary for interoperability. Standards do not dictate *how* the management is accomplished. If an open system wishes to allow a manager to access it remotely, the management protocol provides mechanisms to convey management information.

Management protocol standards provide request/response service between a manager SMAP in one system and an agent SMAP in another system. The agent SMAP carries out management activities on behalf of the manager SMAP. Management standards involve the formal tools necessary for communicating between management entities as well as between layer protocol entities. System management relates to the management of OSI resources and their status across all layers of OSI architecture.

The standards mentioned above do not specify or constrain the implementation entities and interfaces within a computer system. They emphasize only the types of services provided between the station management functions and protocol layers, more specifically the MAC sublayer of the data link layer.

BUS/RING LAN MANAGEMENT METHODOLOGY

We are developing a management application software for IEEE 802.4 Token Bus and IEEE 802.5 Token Ring LANs. Standards 802.4, 802.5, and 802.1 Part B are closely followed with regard to the utilization of management primitives as well as management and protocol entity interactions. Emphasis is on managing the LLC, MAC and physical layers. Primarily, configuration, performance and fault management aspects of token bus and token ring LANs are considered. Centralized and distributed management issues are investigated.

This section and the following sections contain a discussion of various aspects of LAN management research, including methodology, the proposed architecture, and implementation issues.

The management software is being built around a bit-level token bus LAN emulator which has been previously developed.⁸ This emulator implements the MAC portion (interface machine, access control machine, receive machine, transmit machine and the timers) and the LLC portion of the IEEE 802.4 protocol, and is fully operational. The interface management software being built around this emulator relies on the architecture explained in a later section. The manager is implemented as an independent application process (SMAP) resident at each of the stations. The standards are followed as much as possible. The LLC and MAC layers of the current emulator already allow implementation of management primitives. LME part of these layers will be rewritten to reflect the SMAE interaction as shown in the standards.

Both the manager and the agent modes of the management process are being implemented. Stations will be switched into/from these modes by the invocation of SMAP. SMAP will be written and tested as an application process and will be downloaded into each node.

Two managerial modes are developed: centralized and distributed. In the centralized network management mode, only one station's SMAP is switched on and this station is responsible for receiving event reports and sending request and control information. All management information is communicated across the network by SMAE \longleftrightarrow LME \longleftrightarrow PE protocol entity interaction. The manager SMAP in the centrally designated station will also have a video monitor for operator-driven activities and intervention. Network status information will be available to an operator. Operator command information will be an input into a resident SMAP and will be used to either update a SMAP parameter or management command request delivered to other stations through SMAEs.

The second management mode to be considered is the distributed management control. Here, more than one station can act as managers to their designated agent stations. Each manager station uses the same SMAP and operator communication is allowed at these manager stations. The two management operating modes are contrasted in the light of:

- ease of operation and intervention
- complexity of SMAP design

Token Bus/Ring LAN Management Concepts and Architecture

- amount of excess management control traffic generated
- reliability
- recovery from failures
- quality of management service and added-value to network services.

Under these management modes, essentially three issues are explored:

1. configuration management
2. fault management
3. performance management

Initialization, standby monitoring, active monitoring and beaconing procedures are implemented. Configuration changes will be introduced and the manager is asked to manage these through SMAEs. Fault management aspects will be tested by the statistical generation of fault and error conditions. These include station not on ring, multiple tokens, lost token, etc. Manager is expected to catch and rectify these errors.

Configuration Management

Configuration Management involves the initiation, updating and control of information regarding a resource's physical and logical configuration membership relative to the system and to one another. The information normally includes the resource names, addresses, location, resource identification, etc., which are used to monitor and affect the additions, deletions and modifications to network configuration. Configuration input may be the most important activity in most LANs. Configuration parameters are read and set through SMAEs. These parameters must be continually updated especially in more LAN environments where there are a large number of stations and frequent configuration changes.

Performance Management

Performance management is the process of monitoring, tracking and tuning of network parameters to assure a fair, stable and consistent system operation. Performance monitoring is normally passive and includes the following elements:

- packet arrivals and network traffic
- transmissions, retransmissions and checksum errors

- buffer occupancy
- queueing delays and queue sizes
- throughput
- utilization
- token rotation times
- response time
- availability

Performance control is active and aims at correcting, adjusting and tuning the parameters, some of which include the following:

- fairness (to users, stations or applications)
- priority handling
- delay boundedness
- flow and congestion control

Performance tuning is the process of taking direct action by a manager to improve network performance.

Accounting Management

Accounting management deals with the recording of usage charges on a resource basis. These include connect time, cpu, I/O usage and page charges.

Fault Management

In order to be of important value to network operators, the fault management component must offer functions for automatic detection, isolation, notification and, in some cases, correction of failures as they occur. Fault manager also responds to error reports for agent components and may initiate diagnostic tests. Fault management mechanisms are likely to be invoked by event notifications received from layer management entities.

Security Management

Security management consists of functions related to the protection of communication resources and activities. Illegal entry and use of network resources by nonauthorized users are monitored and corrective measures are implemented.

Token Bus/Ring LAN Management Concepts and Architecture

System management implies and presumes a good understanding of the system in question. In order to meaningfully manage the performance of token bus LANs, one must have a very good knowledge of its performance behaviour. There have been numerous papers on token bus and ring LAN performance in which analytic and simulation models have been constructed (for example, References 9, 10, 11, 12, 13, 14, 15, and 16).

Some of the most important parameters which affect token bus LAN performance are the token holding times (THT) and the target rotation times (TRT) for the four access classes. The manager will interact with the MAC layer to control these dynamically, hence preventing starvation for lower class queues and assuring optimum queue depletion rates for each class. There have been a number of studies, mostly simulations but including a few analytical investigations, which have explored the effect of these timer setting on the bus performance.^{17,18,19,20,21} As a result, some insights are now available towards solving the problem. It has been shown¹⁷ that if $THT(i)$ is the token holding time for the i^{th} queue on the bus, and $TRT(i)$ is the target rotation time for the i^{th} queue, then a necessary and sufficient condition for preventing starvation of i^{th} queue is that

$$TRT(i) > \sum_j THT(j) + \sum_j \max(0, TRT(j) - TRT(i))$$

However, this condition can only be used to prevent starvation; it does not assure stability of the queue nor does it ensure optimum utilization of network resources. A preliminary study in the context of file transfers over a token ring²¹ showed us that it is very difficult to achieve maximum throughput for all file transfers. However, sufficient performance tuning is possible by dynamically varying the values of the various timers according to the offered load and the throughput and delay requirements of the applications.

Some of the fault diagnostics and recovery (e.g. lost token) issues are already resolved within the IEEE MAC layer.¹ The manager will address other issues of fault management which are outside the MAC control as described in the standard. These fault management issues and also configuration management will be resolved largely by heuristic reasoning. This reasoning will build upon some of the issues explored by²².

MANAGEMENT ARCHITECTURE

LAN management architecture is designed in such a way as to serve both the central and distributed management requirements. In this study we shall follow

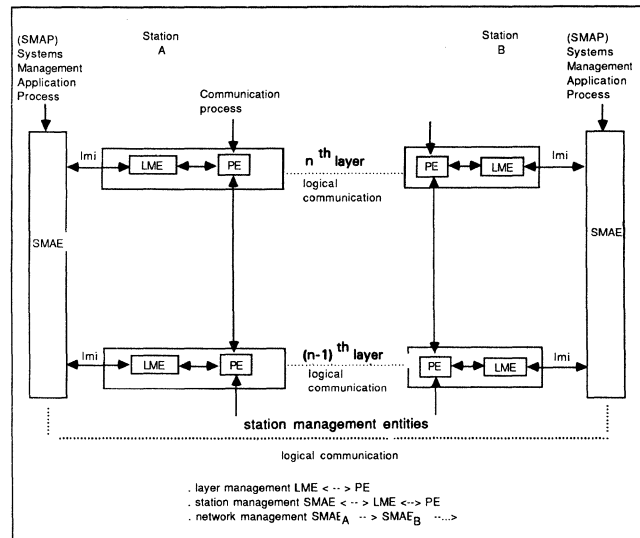


Figure 1. LAN management architecture.

the general outline of an architecture proposed by IEEE 802 standards¹²³. Figure 1 shows the general components of a LAN management architecture at two stations, A and B. The key components are:

SMAP—system management application process

SMAE—system management application entity

LME—layer management entity

PE—protocol entity

Imi—layer management interface primitives

SMAP is the manager/agent process which communicates with the station layers through its SMAE, which surround the “vertical” protocol layers, and communicates with them through the Imis. SMAP may assume roles both as a manager and an agent. It can be set to either mode. As a manager it performs initialization and control of tests and diagnostics and also functions as a collector of reports and notifications. As an agent, the SMAP acts on behalf of the manager. It reports to the manager of status information and fault conditions. It also implements the instructions received from the manager. Some of their lins used in SMAE, LME communications are shown in Figure 4. Layer management is accomplished by LME ↔ PE pair while station management is done by SMAE ↔ LME ↔ PE with the management control logic initiated at SMAP.

For network management logical communication between the station management entities is necessary. Such communication follows the path $SMAE_A \leftrightarrow LME \leftrightarrow PE \dots PE \leftrightarrow LME \leftrightarrow SMAE_B$. In central management mode, only one station’s SMAP acts

Token Bus/Ring LAN Management Concepts and Architecture

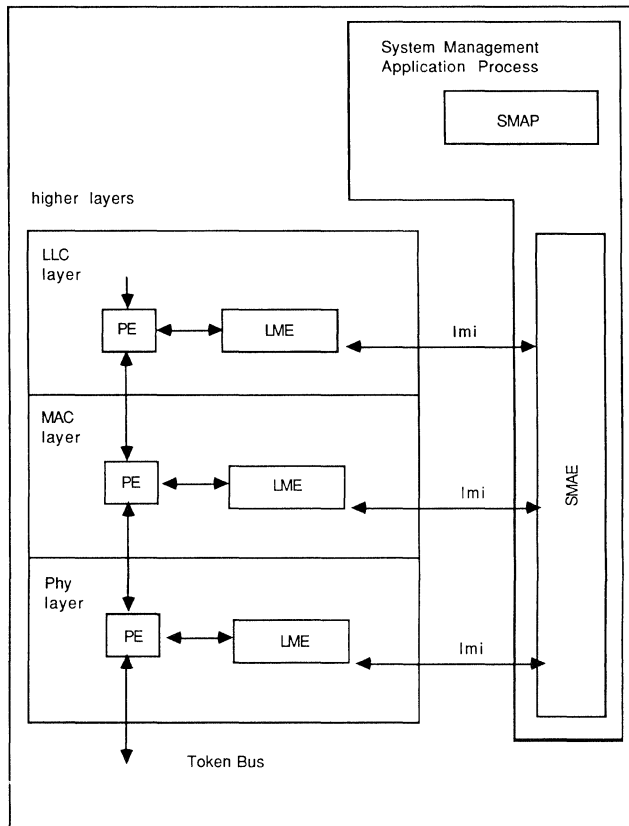


Figure 2. Station view of management process.

as a manager while the other SMAPs are set to agent modes. The system management component is conceptually subdivided into LMEs and SMAE. SMAE is an interface to the management application process. The SMAP is concerned with the management of resources and their configurational, performance and fault status across all layers of the protocol architecture. As mentioned above, SMAE communicates with its own LME's as well as peer SMAEs at other stations. No direct interaction between LME's are permitted at a station. These interactions for interlayer management are accomplished by SMAE. The station and network views of management processes are shown in Figure 2 and Figure 3, respectively. Since most management control for token bus LANs will be affected at physical, MAC and LLC layers, the SMAE interactions are indicated only to these layers.

User processes utilize regular protocol primitives passed from a PE at one layer to a PE at another layer via protocol data units (PDUs) in a conceptual "vertical" flow. The management application process utilizes lmi interface primitives in a "horizontal" flow before they are passed on to PEs. Lmis are normally described in terms of the layer management services provided. The proposed standard lmis, some of




Figure 3. Network view of management process.

which are shown in Figure 4, are defined in detail in the system management section of IEEE Standard 802.4 and 802.5.

DISTRIBUTION OF MANAGEMENT FUNCTIONS

Network management functions, depending on their nature, can be distributed at layer, station and network levels. The decision as to what management

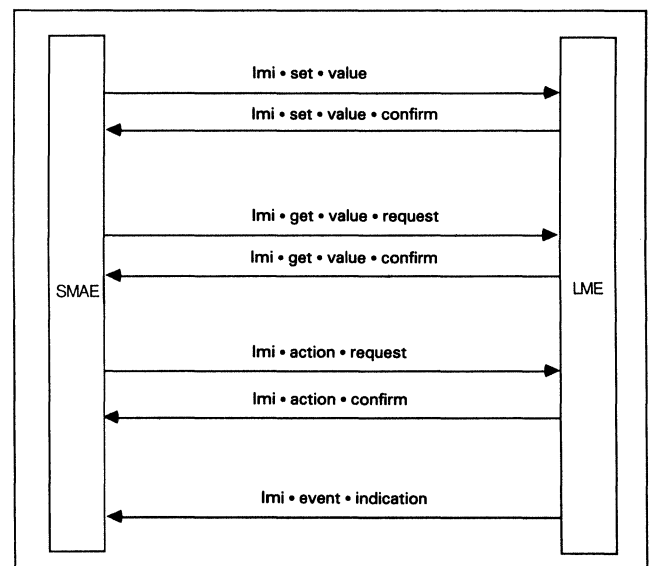


Figure 4. Some lmi primitives used between station and layer management entities.

Token Bus/Ring LAN Management Concepts and Architecture

functionality should be placed at a given layer is a complex one. The following discussion attempts to map a given management function to a location.

Performance Management

In an asynchronous shared media environment, performance control and timing must be based on global objectives. The interdependence of the performance of individual stations precludes an isolated control. Hence performance control is essentially a network level management function. Performance monitoring, on the other hand, may be accomplished at layer and station levels serving as management agents to a network manager.

Configuration Management

Configuration management also seems to be an inherently network level management function. Configuring of topology, software downloading and distribution and topology updating are all network functions. Network manager monitors this information and spreads the data to be used at station levels. Station managers would update their configuration databases to reflect the current topology of network resources.

Fault Management

Monitoring, detection and notification of fault and error conditions are layer and station level management functions but reaction to faults and their correction is a network management function.

Accounting management as it pertains to use of network resources by users connected through stations, is a network level management function.

CENTRALIZED VS DISTRIBUTED MANAGEMENT

As we have discussed in the section on management architecture, the management mode can be configured to reflect either centralized or distributed control. The foregoing discussion on the distribution of management functions serves as a basis of decision and implies that most management functions are resolved at the network level. A protocol is required when the activities between stations need to be coordinated. Indeed SMAP, SMAE discussed earlier provide exactly this protocol. The management architecture and implementation proposed here will

have the following features with respect to centralized-distributed control:

- Any station can be initialized as a network manager
- Several stations can be initialized as managers responsible for certain stations and distributed management functions.
- Stations can be switched on/off into manager/agent modes.

Invoking SMAP with its managerial functions at a station switches it into the manager mode. A manager station communicates with other stations through its own SMAE and peer SMAEs at other stations. This facility assumes a continuity of network management in case the manager station fails. In this case, the management functions are transported to another member of an LAN by activating its SMAP.

Centralized and distributed management can both be implemented. Emphasis will be on the centralized management, however, since most performance, configuration and fault management operations are at a network level rather than a station or layer level.

The network manager station primarily acts as the initiator and controller of the configuration, and collector of station reports and notifications. It requests information from the agent stations and compiles all station reports to derive information with respect to various network status. Based on these, it then dispatches control requests to agent stations. An agent acts on behalf of the manager and reports to the manager all status information about the station including faulty conditions, and also responds to manager requests. It affects and implements manager generated control commands.

IMPLEMENTATION ISSUES

Development of a network management protocol as an application process (SMAP) is a major software undertaking. This process is designed with the following emphasized features in mind:

- both management and agent modes are possible to implement
- configuration management issues, including the initialization and topology update are implemented
- performance management issues, especially
 - queuing delays and sizes for all classes
 - token holding times

Token Bus/Ring LAN Management Concepts and Architecture

- token rotation times
- queue depletion modes
- priority handling and fairness

are addressed, implemented and managed.

- fault management issues including the
 - multiple tokens and lost token
 - nonexistent station
 - multiple station addresses

are implemented.

- comparative merits of central vs distributed management applications are explored.
- relative excess traffic due to management is determined and how critical it is to network performance and stability is studied.
- frequency of manager actions is contrasted with passive monitoring.
- value of operator intervention is examined.

In conclusion, we have proposed a LAN management architecture within the framework of the IEEE 802 standards, and have discussed issues of performance, configuration and fault management. The implementation under way will answer a number of important questions about LAN management and will provide a vehicle to test the merits of various approaches such as centralized vs distributed management. It is hoped that this will help develop formal reasoning processes in LAN management and will pave the way for work on intelligent and expert management systems for future LANs.

REFERENCES

- ¹IEEE Standard 802.4. Token Passing Bus Access Method. December 1984.
- ²IEEE Standard 802.5. Token Ring Access Method. March 1984.
- ³Draft IEEE Standard 802.1 Part B. Systems Management. Revision H. June 1985.
- ⁴Draft IEEE Proposed Standard 802.3. Layer Management. Unapproved Draft. January 1986.
- ⁵Thompson, D.M., "Lan management standards-architectures and protocols," *Proc. Infocom '86, Fifth Annual Joint Conference of the IEEE Computer and Communications Societies*, Miami, FL, 355-363. April 1986.
- ⁶Don Carlos, B., and J. Winkler, "Token-ring local area network management," *Proc. Infocom '86, Fifth Annual Joint Conference of the IEEE Computer and Communications Societies*, Miami, FL, 94-98. April 1986.
- ⁷ISO/TC97/SC21. Information Retrieval, Transfer and Management for OSI. December 1985.
- ⁸Sylvanus, F. and T. Saydam, "IEEE 802.4 token bus emulator." Workshop on Analytic and Simulation Modeling of IEEE 802.4 Token Bus Local Area Networks, National Bureau of Standards, Gaithersburg, MD, 217-228. April 1985.
- ⁹Bux, W., "Local-area subnetworks: A performance comparison,"

IEEE Transactions on Communications, COM-29, 10, 1465-1473. October 1981.

¹⁰Chauhan, V. and A.S. Sethi, "Performance studies of token-based local area network," *Proc. Tenth Annual IEEE Conference on Local Computer Networks*, Minneapolis, MN, 100-107. October 1985.

¹¹Kuehn, P.J., "Multiqueue systems with nonexhaustive cyclic service," *Bell Systems Technical Journal*, 58, 3, 671-699. March 1979.

¹²Rubin, I., and L.F.M. DeMoraes, "Message delay analysis for polling and token multiaccess schemes for local communication networks," *IEEE Journal on Selected Areas in communications*, SAC-1, 5, 935-947. November 1983.

¹³Saydam, T. and A.S. Sethi, "Performance evaluation of voice-data token ring LANs with random priorities," *Proc Infocom '85, Fourth Annual Joint Conference of the IEEE Computer and Communications Societies*, Washington, DC, 326-332. March 1985.

¹⁴Sethi, A.S., and T. Saydam, "Performance analysis of token ring local area networks," *Computer Networks and ISDN Systems* 9, 3, 191-200. March 1985.

¹⁵Sethi, A.S., T. Saydam and V. Chauhan, "An approximate analytic model for token bus and token ring LANs with voice and data traffic." Under preparation.

¹⁶Ulug, M.E., "Comparison of token holding time strategies for a static token passing bus," *Proc. IEEE Computer Networking Symposium*, Gaithersburg, MD, 37-44. December 1984.

¹⁷Nakassis, A., "Token passing networks and starvation issues." Workshop on Analytic and Simulation Modeling of IEEE 802.4 Token Bus Local Area Networks, National Bureau of Standards, Gaithersburg, MD, 102-111. April 1985.

¹⁸Nakassis, A., "On the stability of a token passing network." Workshop on Analytic and Simulation Modeling of IEEE 802.4 Token Bus Local Area Networks, National Bureau of Standards, Gaithersburg, MD, 203-216. April 1985.

¹⁹Muralidhar, K.H., "A hierarchical policy for timer assignments in IEEE 802.4 network." Workshop on Analytic and Simulation Modeling of IEEE 802.4 Token Bus Local Area Networks, National Bureau of Standards, Gaithersburg, MD, 180-202. April 1985.

²⁰Chien, J.Y., "Performance analysis of the 802.4 token bus media access control protocol." Workshop on Analytic and Simulation Modeling of IEEE 802.4 Token Bus Local Area Networks, National Bureau of Standards, Gaithersburg, MD, 102-111. April 1985.

²¹Ananthaswamy, A. and A.S. Sethi, "Some studies of flow control during file transfers in a token ring LAN." Under preparation.

²²Ebihara, Y., K. Ikeda, T. Nakamura, S. Nakatsuka and M. Ishizaki, "Fault diagnosis and automatic reconfiguration for a ring subsystem," *Computer Networks and ISDN Systems*, 10, 2, 97-109. September 1985.

BIBLIOGRAPHY

Coffield, D., and D. Hutchison, "Managing local area networks," *Computer Communications* 8, 5, 240-246. October 1985.

Cole, L.J. "Network management as described in Systems Network Architecture," *Proc. Infocom '86, Fifth Annual Joint Conference of the IEEE Computer and Communications Societies*, Miami, FL, 364-376. April 1986.

LaBelle, N., and K. Chapman, "Building blocks for remote LAN system management," *Proc. Ninth International Fiber Optic Communications and LAN Exposition*. September 1985.

Nakamura, T., K. Ikeda, Y. Ebihara and M. Nishikawa, "Network management in a local computer network," *Software Practice and Experience* 15, 4, 343-358. April 1985.

Sunshine, C., and M. Bernstein, "Requirements for broadband LAN management," *Proc. Phoenix Conference on Computers and Communications*, Phoenix, AZ. March 1985. □

Architectural Support of Network Management: An Alternate View

This report will help you to:

- Discover how LANs can assume a prominent role in their own management, given an adequate management framework.
 - Develop a comprehensive LAN management strategy for your network.
-
-

While it is generally agreed that LAN management involves monitoring, control, and diagnosis, there is no universally accepted definition for LAN management. It means different things to different people, and different things to different organizations. In addition to this lack of definition, other factors make it difficult to develop comprehensive management solutions.

Creative people thrive on innovation. These same people often consider management a mundane task. Historically, LAN system designers did not recognize management as a major design requirement.

The need to share resources is the driving force behind LAN connectivity. For commercial reasons, manufacturers eagerly promote connectivity, fueling the growth of complex systems without regard for their manageability. And since connectivity sells long before manageability, management strategies are constrained by the limitations of installed technology. Connectivity complicates management.

Management is a long-term issue, and *over the long term there is no such thing as a local area network*. The same forces that drive standalone users to connect with LANs also drive LAN users to connect with other heterogeneous networks. Any appropriate manage-

ment strategy must consider the inevitable complex non-LAN environments into which all LANs are evolving. Determining a LAN management strategy from a LAN perspective is planned obsolescence.

A LAN MANAGEMENT FRAMEWORK

The usual approach to LAN management is to identify a need and then develop a utility on a case-by-case basis. It is likely that many commercial utilities were originally conceived by test engineers to help them in their duties, and then embraced by corporate marketing. The problem is not in the utilities themselves, but in the approach. Management support should be built into LAN architectures, rather than derived from a collection of utilities.

The utilities (or tools) of LAN management are line monitors and logic analyzers, management applications, databases, and expert systems. There is good reason for this assortment. Communications is a low-level, real-time function. Information must be available from the lowest layers of communication; hence the use of line monitors and protocol analyzers. LAN management is a data-intensive function. Management facilities must track mountains of data, such as topology, security, and accounting; hence the use of databases. Finally, LAN management is an expert-intensive function. With a shortage of experienced supervisors, expert systems will play an important role in managing complex LAN environments.

This Datapro report is based on "Architectural Support of Network Management: An Alternate View," from a paper given by Dayle S. Woolston and Dale Neibaur, Novell, Inc., at the *IEEE 1988 Network Operations and Management Symposium*, New Orleans, LA, February 28-March 2, 1988. © 1988, IEEE. Reprinted with permission.

Architectural Support of Network Management: An Alternate View

These tools have different relevance in different LANs. Some installations are not interested in security. Others are not interested in accounting. Some may be very interested in monitoring configuration, and not care about fine tuning performance. Still others may be obsessed with performance. The role of applications, databases, and expert systems in LAN management depends on these requirements. For this reason, system designers must build frameworks that allow themselves to be customized.

Figure 1 illustrates a LAN management framework composed of three distributed facilities: a monitor and control facility, a database facility, and a console facility.

The Monitor and Control Facility

The monitor and control facility is distributed throughout all layers of the LAN. Through this facility, the management framework acquires information concerning LAN status. It may also exercise control over LAN components such as initiating diagnostic procedures.

Several constraints govern the monitor and control process. First, information must be gathered continuously. (See Figure 2.) LANs are dynamic systems. Second, information must be gathered from all network components in all functional layers. Managers, operators, and diagnostic applications require timely and accurate information to optimize LAN performance. Third, the monitor and control facility must use network communication resources. It is not practical to have a separate management network to monitor and control a LAN. This separate network would add to the cost of the LAN, and would have to be managed as well. Finally, since the monitor and control facility

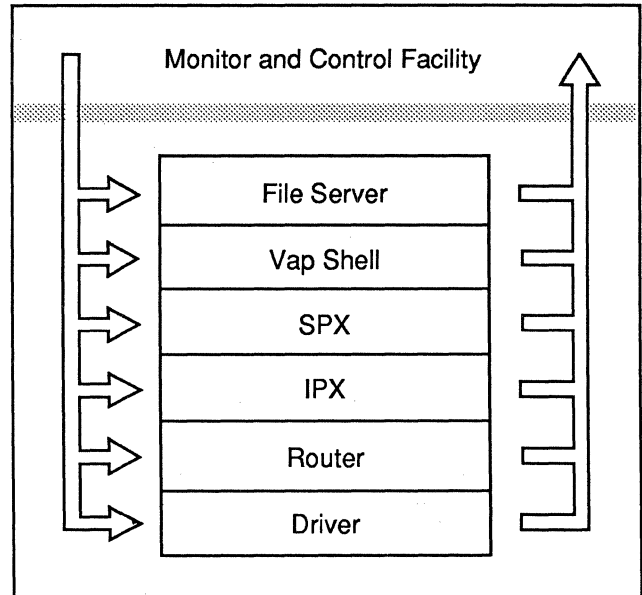


Figure 2. Acquiring information from the LAN—the monitor and control facility.

uses the network, it consumes network resources. The adverse impact of this facility on LAN performance must be minimized.

The roots of the monitor and control facility exist in NetWare LANs, which provide a diagnostic application program interface (API). This diagnostic API provides a window through which monitoring applications may extract information such as driver status, routing tables, and file server connection status. Making such information available is the cornerstone of any effective management framework.

Aspects of Monitoring a Complex LAN

Monitoring the LAN is the most demanding function of the monitor and control facility. Effective implementation requires an understanding of each aspect of monitoring a complex system. The *interactive mode* depends on a dialogue with an operator. Configuration management may require the operator to enter data such as serial numbers and physical location. The *automatic mode* does not require human interaction. High-level management applications may query workstation drivers for transmission statistics.

Two perspectives are involved in monitoring a LAN: *polled* and *spontaneous*. In a polled scenario, a global monitoring process cycles through a list of LAN components (file servers, bridges, workstations, etc.) and requests information. One polled strategy (designated *active*) competes with user traffic. A management application may query file servers for performance statistics at regular intervals without regard for impacting

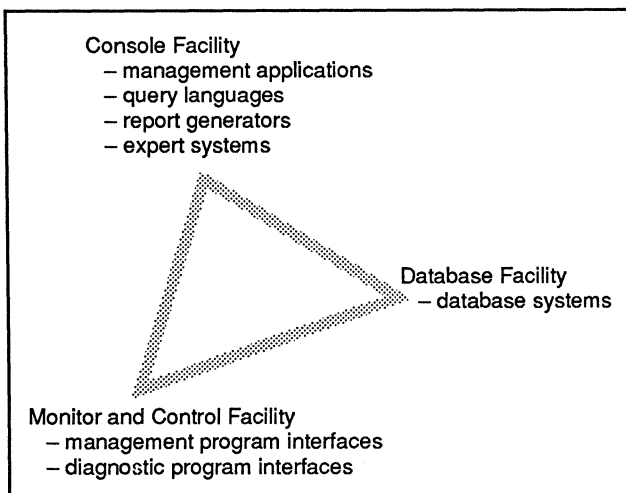


Figure 1. A LAN management framework.

Architectural Support of Network Management: An Alternate View




Figure 3. Monitoring file server status.

user performance. Another polled strategy (designated *passive*) does not compete with user traffic. That same management application may only query late at night when there is no user traffic. In Figure 3, Arrow 1 shows a poll for file server statistics. Arrow 2 in that same figure shows the reply.

In a spontaneous scenario, distributed monitoring processes within each LAN component report to the global monitoring process. One spontaneous strategy (designated *periodic*) means the component reports according to some cycle. For instance, all file servers may report their status every five minutes (See Arrow 3 in Figure 3). Another spontaneous strategy, (designated *event-based*) means the LAN component reports when

a specific event occurs. For instance, file servers may only report their status when they are activated, deactivated, or overburdened.

The Database Facility

In a LAN environment, data can be classified as either *static* or *dynamic*. Static data pertains to LAN configuration (for instance, the type of file server or workstation). Dynamic data pertains to LAN operations (for instance, dynamic routing tables and congestion statistics). A sophisticated monitoring and control facility can provide operators with a flurry of static and dynamic data. The function of the database facility is to store this data.

Several constraints affect the database facility. In particular, management databases must deal with congestion, reduction, and aging. The stream of information headed for storage in the database facility consumes communication bandwidth. The database facility may be distributed to disperse management traffic throughout the LAN.

The flurry of information can also saturate LAN storage capacity. This requires a strategy for reducing the amount of information stored in the database. For instance, one strategy may designate storing minima, maxima, averages, and current values. Another strategy is to change the frequency with which information is acquired by the monitor and control facility. A lower frequency means less information bound for the database facility.

NetWare File Server Console V1.00 Thursday November 26, 1987 10:26 am User SUPERVISOR On File Server NETMANAGE Connection 6			
File Server Statistics Summary			
File Server Uptime: 14 Days 18 Hours 37 Minutes 14 Seconds			
Number of File Service Processes:	10	Current Server Utilization:	5%
Disk Requests Serviced From Cache:	98%	Packets Routed:	0
Total Packets Received:	35,187	File Service Packets:	6
Total Number of Cache Buffers:	49	Dirty Cache Buffers:	0
Total Server Memory:	1,048,576	Unused Server Memory:	5,120
	Maximum	Peak Used	Currently In Use
Routing Buffers:	40	2	0
Open Files:	333	29	25
Indexed Files:	10	0	0
Transactions:	50	40	0
Bindery Objects:	N/A	N/A	N/A
Connections:	100	4	2
Dynamic Memory 1:	16,786	3,230	1,310
Dynamic Memory 2:	35,762	4,632	4,392
Dynamic Memory 3:	16,384	8,798	7,682

Figure 4. A sample management utility in the console facility.

Architectural Support of Network Management: An Alternate View

There is an implicit suspicion of old information in a local area network environment. A file server may be on-line one moment and off-line the next. Aging is a particularly important issue in the case of dynamic data. Outdated information should evaporate from the database facility.

The Console Facility

The console facility is the focal point of the management framework. This facility is composed of a group of management applications that access the monitor and control facility, and the database facility. Operators are thus able to access the data they need to monitor, diagnose, and optimize their LANs. Since management requirements vary widely among complex LANs, the selection of applications in the console facility must also vary. Indeed, this is the reason for defining LAN management in terms of a framework: operators are able to customize management facilities to their particular needs. The selection of applications may include management utilities, database query languages, and expert systems.

The NetWare FConsole utility is a sample component of a LAN management console facility. (See Figure 4.) Such a utility may acquire information from either the monitor and control facility or the database facility. In this case, FConsole acquires information through a special file server API (which may be considered part of the NetWare monitor and control facility). FConsole provides operators with a window into NetWare file server status.

Graphics Applications in the Console Facility. Configuration management is a natural application for computer graphics. An important purpose of a LAN is to share resources. It is common for LAN operators to monitor what resources are available. Graphics applications provide an effective answer to configuration questions through profiling the LAN. Different LAN profiles can be generated depending on whether the operator wants to see all internetwork resources or local network resources only. Profiling can be made even more detailed, based on whether the operator wants to see all available resources, or only those resources currently on-line. (See Figure 5.)

The console facility can make effective use of the database facility by providing a query language interface and a report generator. Operators may then query the database for information such as performance statistics, resource locations, and security profiles.

Analysis in the Console Facility. Diagnosing faults and analyzing performance in a complex LAN requires much expertise. This is a good problem domain for




Figure 5. Profiling network resources in the console facility. The diagram on the left shows those resources available within an internetwork. The diagram on the right shows only those resources available to a certain group.

expert systems. Two types of expert systems can be integrated into the console facility: an expert monitoring interface and an expert query interface.

The monitoring interface provides operators with an expert interface into the monitor and control facility. (See Figure 6.) Such an expert system can be useful in diagnosing LAN faults. The benefits of providing such an interface are many fold. Expert systems do not get sleepy or nervous. They examine each detail of a problem in an orderly manner. Such a tool can be especially effective in guiding less experienced operators.

The query interface provides operators with an expert interface into the database facility. Even experienced operators may find it impossible to identify a problem sifting through a management database saturated with statistics. Such an expert system can be useful in optimizing LAN configuration.

These systems fit into two general categories. An expert monitoring interface provides operators with *con-*

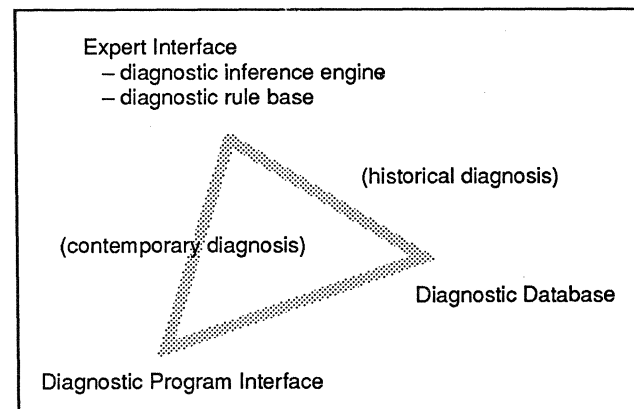


Figure 6. Expert systems in a LAN management framework.

Architectural Support of Network Management: An Alternate View

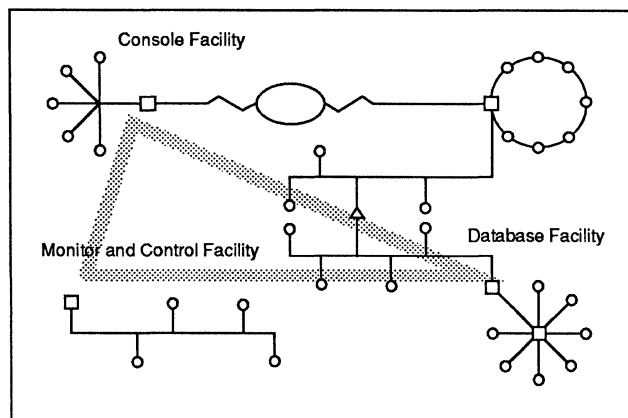


Figure 7. The characteristics of a LAN management framework.

temporary (or near real-time) problem solving assistance. An expert query interface provides operators with *historical* problem solving assistance.

CONCLUSION

The purpose of a framework is to drive management support deep into network architectures and allow customization of management facilities. The management requirements of local area networks cannot be met by a collection of utilities alone. Management frameworks are the first step in defining intelligent LANs. Such LANs assume a prominent role in their own management.

This framework is derived from the fundamental elements of LAN management: monitoring, control, and diagnosis (see Figure 7). The console facility integrates a collection of tools with a reservoir of data acquired by the monitor and control facility, and stored in the database facility. The integration of diverse analytical tools with sophisticated monitoring, control, and storage facilities will greatly enhance the effectiveness of local area network management.

BIBLIOGRAPHY

Chang, D. and S.R. Gross, "Telecommunications Resource Allocation: A Knowledge-Based System", *Proceedings of the Expert Systems in Government Symposium*, The Computer Science Press, October 24-25, 1985, pp. 666-675.

Goyal, S.K. et al., "COMPASS: An Expert System For Telephone Switch Maintenance", *Proceedings of the Expert Systems in Government Symposium*. The Computer Science Press, 1985.

Goyal, S.K., "Expert Systems in Network Management", *Proceedings of the Expert Systems in Government Symposium*, The Computer Science Press, 1985.

Goyal, S.K., and D.S. Prerau, "Expert Systems in Telecommunications—Asia, Americas, Pacific; Pacific Telecommunication Conference Proceedings", 1986.

Guattery, S. and F.J. Villarreal, "NEMESYS: An Expert System for Fighting Congestion in the Long Distance Network", *Proceedings of the Expert Systems in Government Symposium*, The Computer Science Press, 1985.

Guattery, S. and F.J. Villarreal, "A Discussion of NEMESYS Implementation Issues", *Proceedings of the Expert Systems in Government Symposium*, The Computer Science Press, 1985.

International Standards Organization, "Information processing systems—OSI Reference Model—Part 4: Management Framework", *International Standard 7498/4*, Ref. No. ISO/TC 97/SC21N, October 31, 1986.

Missikoff, M., and G. Wiederhold, "Towards a Unified Approach for Expert and Database Systems", *Expert Database Systems*, Proceedings from the First International Workshop, The Benjamin/Cummings Publishing Co., Menlo Park, CA, 1986.

Shepard, Allan and Larry Kerschberg, "Constraint Management in Expert Database Systems" *Expert Database Systems*, Proceedings of the First International Workshop, The Benjamin/Cummings Publishing Company, Menlo Park, CA, 1986, pp. 309-331.

Sykes, E.A. and C.C. White, "Specifications of a Knowledge System for Packet-Switched Data Network Topological Design", *Expert Systems in Government Symposium*, The Computer Science Press, 1985, p. 102-110.

Van Domelen, A.J. and G.B. Bernstein, "Expert Systems in Network Management", *Proceedings of the Expert Systems in Government Symposium*, The Computer Science Press, 1985, pp. 519-521.

Worrest, R., and R. Beretta, *Expert Systems in the Diagnosis of Systems*, GTE Laboratories, 1985. □

Managing Local Area Networks: Fault and Configuration Management

This report will help you to:

- Implement fault and configuration management procedures, which are fast becoming an integral part of successful LAN operations.
- Perform fault diagnosis in both fiber and traditional copper-based LANs.
- Evaluate IBM's approach to LAN management.

LANs are commonly susceptible to malfunctions, such as loose connectors, file server glitches, and application software bugs. Yet only about 10 percent of the organizations now using LANs possess any diagnostics tools—creating a very precarious situation. The cost of LAN diagnostic and management tools is dropping—users contemplating LANs should consider purchasing testing equipment *before* the LAN management issue turns into a corporate liability.

This report discusses the issues, techniques, and tools available for fault and configuration management in local area networks.

THE CHANGING LAN ENVIRONMENT

The use of PCs and workstations increases in the corporate environment as the power of these devices increases. Standalone systems and individual PCs, once considered adequate for most jobs, no longer meet user needs. Users now require connections to their coworkers on a flexible, high-speed network. LANs are becoming the standard intracompany communications apparatus for *Fortune* 1000 firms. Yet, most LAN op-

erating system software lacks the fundamental network management tools needed to monitor users, hardware resources, and data files.¹

During 1988, the worldwide installed base of LANs grew from 500,000 to 800,000. By the end of 1989 that number is expected to exceed 2 million. The worldwide installed base of internetworking equipment, including bridges, routers, and gateways, is forecast to grow from 15,000 units to 600,000 units between 1987 and 1992, reflecting a compound annual growth rate of 60 percent.²

MANAGE THE LAN FIRST, THEN THE NETWORK

A user cannot manage a network without establishing in realtime what is happening on the LAN. When a LAN is first installed, its manager's primary task is to establish connectivity and get the network up and

This report was developed exclusively for Datapro by Daniel Minoli. Mr. Minoli is an adjunct professor at New York University's Information Technology Institute, as well as a full-time data communications researcher and strategic planner.

Index to This Report	Page
Defining LAN Management	402
Fault Management	403
Internetworking Diagnostics	405
Troubleshooting Fiber LANs	407
Configuration and Name Management	408
IBM NetView and LANs	410

Managing Local Area Networks: Fault and Configuration Management

working. Experience shows that the demand for LAN management tools typically lags about one year behind the initial LAN installation. Today's large and complex LANs require sophisticated software-based management and performance-optimization tools. Management must not only prevent deliberate misuse of the LAN but also monitor length and frequency of network access per LAN node.³

LAN managers complain about a lack of quality in LAN network management systems. MIS departments are challenged with accepting the power and flexibility of the PC while retaining their traditional corporate control over software and sensitive data. Today's network management systems were designed to address a communications environment with four primary network components: the (mainframe) computer, carrier-provided circuits, data communications equipment, and terminals. The main function of diagnosis has been to isolate a problem to one of these four general areas and to identify the specific failing element, if possible. These monitoring systems do not work well in a LAN environment.

Fortunately, a new class of software utilities is becoming available to help LAN managers maintain the LAN hardware and software resources. This is a relatively new market which barely existed in 1987.¹ In terms of functionality, LAN management software has come a long way in the past couple of years. Not only are there more products available, but the products are easier to use and much more powerful, with capabilities ranging from audit trail and software metering to system (disk) management, front-ending, and traffic monitoring.⁴

Network management spans a number of key areas. The Open Systems Interconnection (OSI) Network Management standard has defined functional areas pertinent to management activities.⁵ This nomenclature, although originally applied to OSI Management standards, is now becoming commonly accepted for all types of network management, whether OSI based or not. The five categories are described in "Network Management Functions: Telecommunications Hardware," Report NM20-100-101. The following paragraphs relate the methods, techniques, and tools available to the LAN network manager according to these five OSI categories.

DEFINING LAN MANAGEMENT

A pragmatic definition of LAN management is "whatever cures that which ails the LAN: if one can't figure out how to assign user names, network management is assigning user names; if the cable is damaged and one can't figure out where, then network management is figuring out where cables are damaged."⁶

USER SUPPORT TOOLS

Front ends
Document managers
User support
Software meters

SYSTEM SUPPORT TOOLS

User listing
Usage reporting
Performance monitors
Fault recovery

Table 1. A categorization of LAN management tools.

LAN management and software cover a wide range of applications and functions. In broad terms, the LAN management functions can be divided into two classes: user support and system support.

User support capabilities insulate users from the complexities of the LAN and reduce the amount of training and interaction which must be undertaken by the LAN manager. These tools include front ends, document organizers, and remote user support.

System support capabilities give the network manager information about the LAN itself. This aspect of management more closely matches the OSI network management categories. System support utilities provide data on user activity, disk usage, network traffic, system configuration, and so on. These tools complement the lacunas in the network operating system. Some of the available tools address shortcomings of NetWare, NETBIOS, and MS-Net operating systems. In terms of products, LAN management tools can be grouped in the areas shown in Table 1. In the following paragraphs, these tools are mapped to the five industry-accepted classes shown in Table 1, although in a few cases, the mapping is a matter of judgment (particularly for the user support tools).

Most workgroups probably started as a few PCs sharing resources but, with increased processing power, many have expanded to include thousands of PCs connected with hundreds of servers and other shared resources. Multiple software protocols and interfaces will become more commonly included in a single server or gateway interface. Network management will be very important in this setting. A network manager preferably wants to monitor the entire network with one tool and to view it as one consolidated entity. At present, there is a diversity of vendors, interfaces, protocols, and devices, making network control and management an ever-changing challenge.

Network managers may be forced to use several tools to get a complete picture if the network happens to contain a mix of IBM equipment, non-IBM main-

Managing Local Area Networks: Fault and Configuration Management

frames or minis running TCP/IP, gateways and bridges, and workgroups that share files and printers. In this scenario, consolidating the information onto one display is difficult. Microsoft's OS/2 LAN Manager represents a de facto workgroup networking standard that includes extensive network management support. It has gained rapid acceptance among network vendors and applications developers.

Network managers are increasingly witnessing an erosion of their control over the networks they must support. This growing trend makes the task of network management even more difficult. The increased usage of decision support systems in the LANs has left many managers with little or no control over what is plugged into their networks. The following items comprise a wish list of desirable network management features:⁷

- Expert systems for diagnosis of (common) LAN problems;
- A mechanism whereby LAN alerts for the network administrator are generated automatically when problems arise;
- Online LAN help facilities; and
- A central LAN backup system.

Automatic Alerts. The Microsoft OS/2 LAN Manager issues automatic alerts when a security and/or performance problem arises. LAN Manager can also alert the user directly for problems that may be under his/her control (for example, "printer out of paper"). Novell NetWare tracks system statistics that indicate the status of the LAN and a number of LAN services; however, problem notifications are not proliferated.

Online Help Facilities. Users should have the capability to query a LAN server, which can provide static answers to the most-asked technical questions. An enhancement would be an expert system that leads the unsophisticated user to a solution, in easy-to-follow steps. The file could contain product name and configuration information such as interrupt usage, input/output address, and port location. Installation software for new products could check the file to self-configure or recommend a configuration.⁷

A Central LAN Backup System. Users need a secure and reliable way to perform file backup. The LAN administrator must have read/write access to every file connected to servers and PCs on the LAN. To minimize the LAN's vulnerability to unauthorized access, backup services must come from the LAN vendors. Interworking between the LAN operating system and backup software should include the capability to encrypt data files based on the identity of the owner.

FAULT MANAGEMENT

Fault Management is the discipline of detecting, diagnosing, bypassing, repairing, and reporting on network equipment and service failures. There are three basic steps in LAN fault management and troubleshooting.⁸

Step One. Understand the particular LAN you are managing. What are its characteristics? Is it a ring, bus, or star? What type of cable is used? How is the cable laid out, and how long are the runs? How many servers are there, and how are they configured? What is the network operating system and version? What level of DOS is installed, and is it the same for all workstations? What types of applications are users running?

Step Two. Establish a logical procedure for sectionalizing the problem. One can often pinpoint the problem by moving, replacing, and testing the cables, servers, and workstations.

Step Three. Apply the proper tools to diagnose problems that are not immediately obvious.

Users typically tackle fault management by employing hardware tools. LANS can break down in three areas: hardware, software, and cabling. Of the three, cabling is the most common problem, based on empirical observation. For the purposes of this report, cabling includes both the cable proper and the interface card.

CABLING

Cards. Network interface cards, cables, and other low-level hardware all have an impact on the operation, performance, speed and throughput of the LAN. If interface cards and cables are working properly, one cannot expect speeds greater than the rated throughput. If an interface card is working improperly or a cable is not correctly terminated, errors may occur. In this case, the LAN will run significantly slower due to re-transmissions of mutilated and lost packets. A defective card can significantly slow down a network.

There are a number of ways to check if interface cards and cables are working properly. Most LAN manufacturers include a basic diagnostic utility program integrated with their hardware. While these programs can detect severe errors, they may not show subtle errors. These utilities may report that they see other workstations on the LAN and that, therefore, the LAN is "working" although a real problem may exist.⁹

Some available products monitor traffic on the LAN and indicate trouble spots and can be used to test the cable system. These diagnostic programs will generally provide much better information about the physical LAN than the simple diagnostics included with the in-

Managing Local Area Networks: Fault and Configuration Management

terface cards. There are also cases where a mixture of different vendors' Ethernet cards may not work together. Diagnostic programs included with vendors' cards may help indicate any incompatibilities; with the use of more sophisticated hardware diagnostic programs, an experienced installer should easily be able to determine compatible card mixes.

Once the data is read from the server, it must be transmitted over the LAN, to the station which requested it. Unless the network itself is optimized, there may be no benefit to having a fast server; hence Performance Management and Fault Management must be done in a cooperative fashion.

Typical fault diagnosis products evaluate the network's capacity, providing traffic monitor functions by reporting network bandwidth usage (in percent). In addition, these products detect physical problems on the network, such as impaired cables, hubs, and repeaters, as well as bad network interface cards. The cable test generates packets, then measures the success rate of the transmissions.

Some fault diagnosis products can indicate how many packets collided and the number of bits lost or dropped; certain products also report the number of reconfigurations necessitated by lost tokens. Network managers can use this information to formulate fault scenarios.

Wiring. A network manager can perform continuity, loopback, and reflectometry tests to diagnose cabling problems. (Cable installation companies may also perform reflectometry tests on behalf of the manager.) These tests help ensure that the cable system is working correctly. The most common problem with LAN cabling is broken cable or improper termination. Most connectors do not handle movement well and, if left lying on the floor, they may be stepped on and broken. A regular inspection of the cable ends and connectors will help identify defective terminations. There is a wide range of tools available for diagnosing cable problems.¹⁰

Ohmmeters. An ohmmeter is a simple tool that gives impedance measurement. One can use an ohmmeter to locate open or shorted cable. If the impedance reading matches the rated impedance of the cable, the cable is fine. If the reading does not match the rated impedance, then the LAN has a short, a crushed cable, or a cable break somewhere along the cable run. Ohmmeters cannot be used in fiber-based LANs. Ohmmeters are typically priced well under \$100.

Outlet Testers. Sometimes the problem is not the cable but rather the electrical outlet. For example, if an outlet is not grounded properly, noise or even current may

be introduced through the power supply into the workstation, and then through the network interface card onto the copper-based (twisted pair or coaxial) LAN cable. An outlet tester, which costs around \$10, can help detect this type of problem.

Coax Connectors, T-Connectors, and Terminators. Extra terminators are a basic requirement, particularly on a bus network. Extra terminators can be used to isolate sections of the cable for testing.

Oscilloscope. An oscilloscope allows the network manager to examine the cable's waveform. An oscilloscope helps detect the existence of noise or other disturbances on the wire, such as continuous voltage spikes. Again, this applies only to copper-based LANs.

Time Domain Reflectometers (TDRs). A time domain reflectometer operates by sending an electrical pulse over the LAN cable, monitoring for signal reflections. On a good cable there will be no reflections, indicating that the cable is clean, with no breaks or shorts. If there is a break or short in the cable, however, the time it takes for the pulse reflection to return gives the TDR a very accurate idea of where the fault is located. Many TDRs can locate cable breaks to within a few feet. TDRs have traditionally been relatively expensive instruments; a TDR with an oscilloscope costs \$5,000 or more. A new, less expensive generation of TDR equipment is now available which costs \$1,000 or less. These new instruments are more compact than their predecessors, often measuring about the size of a paperback book or smaller. The newer TDRs are also easier to use and still accurate to within a few feet.

FIBER OPTIC LAN DIAGNOSTIC TOOLS

Fiber-based LANs require different equipment. While they provide significant advantages over conventional LANs, the fiber LAN's design necessitates more sophisticated test equipment. Fiber LAN diagnostic tools must provide comprehensive design verification, including the capability to precisely determine bandwidth, sensitivity, and linearity. These measurements can only be performed with fiber optic test equipment that includes sophisticated parametric capabilities.

The market is experiencing an influx of test equipment that helps technicians diagnose and maintain a fiber-based wiring system. These instruments vary in complexity from pocket-sized power measuring units to console-type testers. Only those instruments that are optical in nature can make performance measurements on the optical parts of any electro-optical system. Fortunately for electronics test engineers, the few optical instruments used outside research laboratories are relatively simple in design.

Managing Local Area Networks: Fault and Configuration Management

Fiber optic instruments can be divided into three general categories: 1) Power Meters (optical loss test sets); 2) Optical Time Domain Reflectometers (OTDRs); and 3) Optical Bandwidth Test Sets (OBTSS).

Power Meters. Power meters (optical loss test sets) measure the optical power from a length of fiber in much the same way that conventional power meters measure electrical power. These tools are used to perform a one-way loss measurement. The loss may occur in the fiber, connectors, splices, jumper cables, and other system areas. Some power meters also have a built-in transmit source. Two sets, both with transmit and receive capability, are used together to make measurements in both directions without having to relocate personnel or equipment. The individual power meter is a single unit consisting of an optical receiver and an analog or digital readout. A light source, typically at the point of origination, supplies the power which is detected at the end of the fiber link or test access point. The meter displays the power detected in decibels. The wavelength range often given indicates the various wavelengths that the meter can detect. The resolution parameter indicates the smallest step that the meter will display.

Optical Time Domain Reflectometers (OTDRs). Network managers can use OTDRs to characterize a fiber wherein an optical pulse is transmitted through the fiber and the resulting light scattered and reflected back to the input is measured as a function of time.

OTDRs are useful in estimating the attenuation coefficient as a function of distance and in identifying defects and other losses. These devices operate on basically the same principles as a copper-based TDR. The difference is one of cost: OTDRs typically range from \$10,000 to \$18,000.

Optical Bandwidth Test Set. Optical bandwidth test sets consist of two separate parts: the *source*, whose output data rate varies according to the frequency of input current applied to the source (specified by frequency range parameter); and the *detector*, which reads the changing signal, determines the frequency response, and then displays a bandwidth measurement. The instrumentation is calibrated using a test fiber; the actual measurement results are compared to the calibrated value to display the bandwidth value.

HARDWARE

Until recently, LAN hardware troubleshooting was limited to built-in LAN diagnostics. Most LAN troubleshooters still use manual, step-by-step procedures

for diagnostics. While there are many tools for troubleshooting Ethernet networks, fewer are available for token-ring LANs.

Ethernet LANs can be checked by sending traffic to a particular node and waiting for a response; one can compare that response to the response of another node. A fault exists if the node in question is sending packets and receiving none in return.

The network manager can also troubleshoot an Ethernet LAN by counting or sampling packets on a network. In this manner one can, for example, determine how much traffic a given server is handling. One typical problem is loss of information, i.e., mutilated or incomplete packets. Packets may fail to reach their destination intact if the cable has been extended beyond the vendor-recommended length.

Many networks can perform loopback tests in which a workstation sends a message to itself. If the packet is not received intact then there is a problem, usually with the network interface card.

Recently, some vendors have introduced a number of products that compensate for the simple diagnostics included by the manufacturer in the network hardware. Users can install this additional equipment either as a network node (operating in monitoring mode) or connect it only when a specific problem exists (testing mode). In monitoring mode, the equipment checks network parameters such as preamble length, alignment and CRC faults, frame lengths, and collision rates. In testing mode, the equipment allows the troubleshooter to run interactive diagnostics set on different parameters, notifying the troubleshooter when a given level exceeds a given threshold. This new hardware monitoring equipment ranges in price from \$300 to \$5,000, depending on sophistication.

INTERNETWORKING DIAGNOSTICS

As internetworking becomes commonplace, diagnostics problems grow more critical. Troubleshooting across gateways and bridge links is a real challenge. Although LAN vendors have not yet been able to offer comprehensive network management products for a single LAN, users are already demanding more sophisticated products that can manage integrated multiple LANs over a geographically dispersed area, such as a Metropolitan Area Network (MAN) configuration. Network management limitations continue to be the most frequent reason why users limit the size and scope of LAN implementations, according to experts.³

A whole new set of tools is necessary to monitor the traffic between networks. Some vendors are already

Managing Local Area Networks: Fault and Configuration Management

building these tools into gateways and bridges.¹⁰ Currently, fault management revolves around diagnostic tools and software used by network administrators to monitor and diagnose hardware. The next logical step is to maintain and diagnose networks through expert systems.¹¹

While network managers wait for more sophisticated tools, they can employ the following suggestions to troubleshoot problems for effective fault management⁸:

- **Technical Bulletins.** Current technical bulletins and bug fixes can save hours of testing and frustration. End users should request copies from the dealer.
- **A Reliable Interface Card.** If the benchmark board card is not trustworthy, then all results will be inconclusive when testing a suspect station by switching the interface card.
- **Walkie talkies** can save time, particularly when servicing large, spread-out installations.

TROUBLESHOOTING SOFTWARE WITH PROTOCOL ANALYZERS

Protocol analyzers are important high-end tools for fault management. The analyzer is a specialized workstation that collects, analyzes, and displays the data circulating in a LAN cable. An operator using an analyzer literally sees everything on the cable: PDUs, messages, files, passwords, IDs. An analyzer can display the information and store it to disk for future study.¹²

Protocol analyzers assist in diagnosing Ethernet-based LANs by breaking out information on protocols ranging from TCP/IP, OSI, and DECnet to IPX, AppleTalk, and NFS. Protocol analyzers can solve a wide variety of network troubleshooting problems; however, a fair level of technical expertise is required to assess the results. Additionally, protocol analyzers are expensive, ranging from \$17,000 to \$30,000 depending on the number and types of protocols supported. Both token-ring and Ethernet models are available. While dozen of vendors build traditional protocol analyzers for X.25 and BSC/SDLC networks, only a handful of vendors currently build LAN analyzers.

Speed is a major problem with network analyzers, however. Protocol analyzers are not designed for full packet capture in a heavily loaded environment. Most analyzers barely keep up with network utilizations of 30 or 40 percent; very few can handle 60 to 70 percent utilization on Ethernet. To be truly effective, network analyzers must provide expanded packet buffers; some models now feature buffers of 8M bytes or more.

The Open Systems Interconnection (OSI) and MS-Net protocol suite now includes full OSI Session and OSI Presentation Layers. Consequently, sophisticated analyzers must support layers one through six of the OSI model.

Typically, protocol analyzers are primarily used in one of two circumstances:

- 1) To monitor the network to identify why performance is degraded.
- 2) To resolve explicit problems detected by other network monitoring or management software.

Protocol analyzers support the technique of "stepwise refinement," which is a pragmatic course of action in troubleshooting networks. A five-point outline which uses step-wise refinement is listed below.¹⁴

1) Begin with a generic test to observe the current status of the network. A protocol analyzer works well here, since it collects all packets on one channel and several types of errors on other channels.

2) Look for high utilization. 60 percent or higher is considered the region where Ethernet performance begins to degrade. Also look for high error rates.

3) Try to isolate high traffic- or high error-producing stations on the network (very often, they will be the same).

4) Look for high counts of broadcast or multicast packets. In internetwork configurations with bridge, router, or gateway problems, large numbers of broadcast or multicast packets can slow down network performance with spurious traffic.

5) Look for high CRC/alignment or short packet counts, indicating media or connection problems. In general, it is advisable to eliminate physical media-related factors before proceeding to investigate software as a potential source of network problems. The network manager should verify that name tables and/or name servers are defined correctly and working properly. Also, he/she should ensure that multiple nodes are not contending for the same Ethernet address. In summary, network managers should eliminate the obvious problems first before shifting to more sophisticated protocol analysis. Empirical observation shows that 90 percent of the fault problems can be identified by using straightforward step-by-step methods such as the one just listed.

If these steps fail to isolate the problem, then one must tackle the protocols. Protocol analysis can be particularly helpful in identifying whether the application or

Managing Local Area Networks: Fault and Configuration Management

the transport mechanisms are causing problems. It is advisable to start from lower levels of the protocol stack and work up the hierarchy, thereby isolating the problematic layer through the process of elimination. To accomplish this, the network manager must collect the PDUs for the various layers and components of the protocol, beginning with the lower layers of Ethernet and proceeding all the way to the application environment. Most protocol problems occur at the interfaces between layers or between software systems. Typical problems are values that are out of range or out of order, too long or too short fields, missing or incomplete data, and garbled information due to incorrect reformatting.¹⁴

Some fault (and performance) management applications of protocol analyzers follows.

If, for example, the network manager suspects a high number of erroneous or undelivered packets, he/she may use a Names Table constructed for the network in conjunction with a protocol analyzer to monitor all traffic and observe receive and transmit error counts for the network stations. Those nodes with high transmit and/or receive error are good candidates for further investigation.

If uneven traffic distribution on the network appears to be impacting some or all users, it is often useful to segment the network and separate networking workgroups using bridges that can filter extraneous traffic. The protocol analyzer can assist in this segmentation process by showing which nodes communicate with each other most frequently. This can be achieved by building channels that trap all server-related packets on a channel-per-server basis, and then examining the trace to see which stations access which server. As long as servers and heavy users are kept together, segmentation should improve overall network performance. If the network manager wishes to determine the impact that adding new servers and users will have on existing network users, he/she can use a protocol analyzer to simulate the load produced by the new network devices and monitor overall performance.¹⁴

The extensive capabilities of analyzers make them a potential security threat, however, allowing a hostile user access to anything that runs on the network. LAN cables are "party lines": every frame reaches every station on the LAN (although in bridged LANs, frames do not usually leave the local LAN for which they are destined). By agreement, each LAN workstation looks only at frames addressed specifically to it (as well as at broadcast frames, which are addressed to everyone, and at frames sent to group addresses). Protocol analyzers, however, violate this agreement and read every frame or an operator-specified subset of frames.

TROUBLESHOOTING FIBER LANs

An optical LAN connects devices using optical fiber (typically multimode). Although optical LANs are available in three distinct topologies—star, bus, and ring—an estimated 80 percent of all optical LANs today use ring topology.¹⁵ About 95 percent of the fiber LANs use graded-index fiber; the balance is step-index and plastic clad fiber. The use of single mode fiber for LANs is still very low.

The transmission speed of optical LANs is increasing each year. Recently, 10M to 50M bps speeds have gained attention; vendors have also introduced backbone optical LANs with transmission speeds of 100M bps and higher. This is no major technological breakthrough, however, since today's fiber networks can easily carry 2.4G bps for field-deployed technology, and higher in the laboratory.

Optical loss is an important item to measure when troubleshooting fiber LANs. A stabilized light source and optical power meter are key elements of this process. Each optical fiber and circuit part used in optical-fiber communication produces an optical loss, which is one of the factors that limits the range of the optical communication system.

A stabilized light source, a mode scrambler, and an optical power meter are used to measure this optical loss. The optical-output stability of the stabilized light source is important for increasing optical-loss measurement precision.

Two devices are used as the light-emitting element of a stabilized light source: LDs and LEDs. Each has unique properties for measurement purposes. The LD has a high output level, narrow light emission spectral width, large optical-output change versus ambient temperature change, and high coherence. The LED, on the other hand, has a low output level, wide spectral width, small optical-output change versus ambient temperature change, and low coherence.

Because of these different properties, the LD stabilized light source is used in high-loss measurements and the LED stabilized light source is used in low-loss measurements that demand greater precision.¹⁵ The demands for an optical-power meter differ between high-capacity, long-haul communications systems and optical LANs. The repeating zone is longer in low-loss optical fiber used for long-distance communications; thus, the optical-power measurements are made over a wide dynamic range. This is not true for optical LANs, where the requirement is for simple operation.

Managing Local Area Networks: Fault and Configuration Management

FAULT RECOVERY

Protecting data against hardware failure is an important issue, which can fall under the auspices of either fault management or security management. Reliable software designs that ensure against data loss in the event of hardware failure are becoming more prevalent in the LAN market. Several methods of fault recovery are currently available.¹⁶

Disk Bad Track Handling. Very few disks have no flaws; thus, it is important to provide software that can detect flaws and deal with them transparently.

Mirrored Disks. This involves writing data to two separate disk drives so that both drives will contain the same data. In the event that one drive has an error, the alternate drive will continue operating without interruption to the LAN system. For example, Novell's NetWare offers duplexed drives, each containing two drives and separate controllers. While this is an effective approach to protection from disk failure, it does introduce a higher hardware cost.

Transaction Commit, Concurrency, and Recovery. This feature gives an application the capability to protect data files from application failure. By grouping several I/O requests into a single transaction, the operating system will not write the transaction to disk until the application has terminated or issues a commit command. Novell's NetWare and Univision's LifeNet offer this capability.

Transaction Logging. Transaction logging, which is essentially realtime backup, is a powerful fault management feature. After data is written to disk, the information is "echoed" to the server's local tape unit. As each I/O is performed on files, the data is written to the tape drive. If there is a system failure, the file state may be recovered by repeating the I/O requests for that file from data written to the tape. Currently this capability is unique to LifeNet.

CONFIGURATION AND NAME MANAGEMENT

Configuration and name management is concerned with maintaining an accurate inventory of hardware, software, and circuits as well as the ability to change that inventory in a smooth and reliable manner in response to changing service requirements. Configuration management affects network design, performance issues, and even security.

One of the most basic issues for the LAN administrator is maintaining system configuration maps. System configuration is the list of system parameters showing

who has access to given network software and given databases. While most LAN operating system software gives the administrator the ability to add, delete, and modify system configuration parameters, it typically provides little functionality in monitoring system configuration.

One configuration management technique to improve network performance is to use a bridge or router to break a single network into two separate but logically connected LANs.

Each time another user is added to a network, more workstations share the bandwidth. Eventually, if the network is heavily used, that bandwidth becomes saturated and network performance suffers. Bridges and routers are a solution to this problem. Both bridges and routers link two LAN systems into a single system. The two separate LANs can still communicate transparently, but only internetwork traffic (traffic intended for the LAN on the far side of the bridge) actually passes over the bridge and circulates in that network's cabling system. Each LAN's local traffic remains local and each network will use up about half the bandwidth—a contrast to the situation where the two LANs are combined into a single network.

There are a number of bridge and router strategies in use today. Bridges at the Media Access Control (MAC) layer are usually external devices that work transparently to any LAN software. MAC layers usually link similar hardware systems; the most common are Ethernet-to-Ethernet bridges. Routers are protocol dependent, hence they only work with one network operating system. Under these network operating systems, several different network cards can be plugged into the bus of the file server at the same time. These different interface cards each run two separate networks; the file server software will link them so that the different networks can still communicate. Routers internal to file servers are popular in PC LANs. A NetWare server, for example, can internally bridge up to four separate networks. NetWare file servers' internal routing can also be used to bridge dissimilar hardware systems.

LOGICAL ARCHITECTURES

While there are many types of cabling and network interface hardware, there are only two basic logical architectures that are currently supported on PC LANs: 1) peer to peer and 2) client/server.

Peer-to-peer architectures require no dedicated file server, because any node on the network may share its local hard disk with other nodes on the network. This type of architecture is used in IBM's PC LAN Program and 10net Communications' 10Net. Peer-to-peer architectures are often used in smaller LAN installations,

Managing Local Area Networks: Fault and Configuration Management

as they require no additional hardware and generally have a lower cost per node. Peer-to-peer architectures may create security problems, however, due to the lack of centralized data storage (which is more easily physically and logically protected). Peer-to-peer LAN networks also have slower performance and require greater administrative effort to configure and maintain security definitions.

The client/server architecture depends on services provided by dedicated file and print servers and thus requires additional hardware. The higher cost of client/server LANS is often compensated for by the higher performance and more reliable security. The centralized disk storage architecture provides a mechanism for controlling both user access and backup operations.

When determining network operating system (NOS) configuration, the network manager should classify security needs into three categories: 1) minimal or no access control; 2) medium access control; and 3) maximum access control.

For installations requiring little or no security control, any NOS on the market can be adequate. Lowest cost implementations can be configured by use of peer-to-peer LAN software.

Higher security requirements will require selection of a NOS with more specific and sophisticated security. See Table 2.

MINIMUM SECURITY LAN
Peer-to-peer architecture
DOS disk format
Bootable workstations (local storage)
No directory or file access control
Shareable printers cross the network
MAXIMUM SECURITY LAN
Dedicated file server
Non-MS-DOS disk format
Diskless workstations (remote boot)
Access control down to lowest level (file)
Password encryption
Security monitoring and accounting
Network encryption devices
Printers attached to secured file server
Automatic logout
No remote login
Reduced system privileges
Fault-tolerant design

Table 2. LAN configurations and their security requirements¹⁶.

DOCUMENT MANAGERS

Document managers help end users locate files and documents and require neither the full DOS name nor the exact location of the directory. With document managers, users need to know only basic information to locate a file; users can call up a file by author, subject, project name, or creation date.¹⁷ Users can search on any field to find a file anywhere in the network.

Several document managers also load the software that created the file, allowing the recipient user to work on the document. Other document managers can work in conjunction with E-Mail packages.

Document/file organizers are very useful when many people need to access a document. These systems also act as front ends, since they insulate the users from the complexities of the directory structure. Document/file organizers range in price from \$500 to \$1,000.

REALTIME USER TRACKING PRODUCTS

The first step in collecting network information is to determine who is using the network as a function of time. In instances when all users must log off or be logged off (for example, before doing a backup), this type of tool is useful.

Realtime user tracking tools provide value-added functions that go beyond the listing of active users generally available from the network operating system. Some of these value-added features include graphical floor plans, identification beyond bridges, and sorting capabilities. Prices range from \$100 to \$300.

END-USER SUPPORT

Many problems which end users attribute to LANs are, in fact, problems that stem from the user's lack of familiarity with the various LAN applications he/she wants to employ. Network managers can resolve such problems by using remote access tools to determine the cause of the application problem.

Remote user-support tools allow the network manager to gain access to the remote workstation or PC; access to the keyboard and screen can be acquired remotely so that the manager can determine the possible problem source. In some systems, the manager is automatically informed of the user's hardware. In addition, some of these systems maintain statistics on who was helped, the nature of the assistance, and the duration of the session. This data can be used to bill back the support costs and/or to design training programs. Prices for these systems range from \$300 to \$1,500.

Managing Local Area Networks: Fault and Configuration Management

FRONT ENDS

Front ends are menu-driven interfaces that insulate the users from the native environment. The menus will convert from the user-friendly commands to the necessary DOS or LAN operating system commands. End users should not be expected to understand DOS, the LAN software protocol stack, and other details such as software version numbers.

Many front-end systems provide a template that the network manager can customize for different LAN user groups; these menus can typically be produced at the manager's terminal and distributed over the network. A front end may perform additional functions as well. Some front ends provide basic usage tracking. While these reports are not as detailed as those produced by sophisticated audit trail tools, front-end systems can report how much time a user spends in a particular application. Others may blank out the screen after a period of inactivity or even log off the user. Other front ends have a software meter, which monitors applications usage, to support consistency with the license restrictions. Front-end systems range in price from \$100 to \$700 per server.

IBM NETVIEW AND LANs

IBM's NetView allows users to identify and correct problems in a traditional SNA teleprocessing network. Announced in 1986, NetView allows key nodes of a large SNA network to send alarms and alerts to the IBM mainframe. In addition to control and problem determination features, NetView options include NetView Call Accounting (developed by DMW Communications, Inc.), which charges back voice system cost, and NetView Voice Network Design (developed by Vector Software, Inc.) to aid in designing voice networks.

Additionally, IBM offers NetView/PC, which acts as an interface to NetView for non-SNA devices. The interface does not provide any functionality by itself: a LAN manager program must generate the information; NetView/PC formats it and sends it to NetView. Technically, NetView has the potential to carry out centralized network management for a large collection of local workgroup LANs; however, the problem is one of jurisdiction.

In spite of NetView's lack of sophisticated support for non-IBM LANs and other non-SNA elements, there are no comprehensive alternatives today—nor will there be for at least two or three years. NetView has a two-year head start on AT&T's Unified Network Management Architecture (UNMA), and over three years' jump on Digital's Enterprise Management Architecture (EMA). According to industry watchers, if real

standards do not come out soon, NetView will become a solid de facto standard.⁶ The Institute of Electrical and Electronic Engineers (IEEE) is still working on the 802.1 LAN management standard.

As of 1988, no major LAN vendor had a working NetView-compatible product. 3Com has been testing NetView.⁶ A number of vendors, including Novell, have announced NetView interfaces for their products. Most networking vendors will probably feel obliged to support NetView to some extent.

NetView has a centralized architecture, however, and many LAN network managers prefer a more decentralized approach. Some industry watchers guess that the big LAN vendors will try to develop superior, low-cost proprietary management systems for their LANs, while providing gateways to NetView.

Architecturally, IBM's network management approach separates the management of a network from the network's data transport. Within this framework, network components acquire two types of responsibilities: transport and management.¹⁸ The IBM approach defines four entities within the management framework: 1) focal point; 2) service point; 3) entry point; and 4) target. These entities interact to take advantage of the capabilities that each network component offers to the rest of the network.

A **focal point** offers the functions required to manage the network from a central location. The focal point manages all of the remotely and locally attached network components. This would be NetView, in IBM's product parlance.

The **service point** makes a network component visible to the network management system residing at a focal point. The service point handles the interaction between the focal point and the network component, especially with regard to transporting management information. This is NetView/PC, in IBM's product context.

The **entry point** is a network component from the perspective of transporting information through the network. It assumes the responsibility of a service point, providing network management information about itself to the network manager.

The **target** is a network component that does not have its own direct access to the network management system. It is a device or subsystem that provides only information transport; the management capabilities are provided to it by a service point.

The focal point/service point/target approach is applicable to all network components regardless of their

Managing Local Area Networks: Fault and Configuration Management

structures or transport characteristics; however, this approach is particularly appropriate when the LAN subsystem is considered a target and the management server of the LAN is coupled with the functions of a service point. IBM's LAN Manager provides the LAN-specific operation at a service point, reporting to a centralized focal point. The NetView/PC program provides the service point function which is linked to NetView; the focal point is in a mainframe. In this fashion, the hierarchy of network management is extended into the LAN.

IBM's LAN Manager can provide centrally accessible control for the management servers distributed to the rings through logical link connections, with the LAN Reporting Mechanism residing on each server. Statistical or problem information can be forwarded to the LAN Manager; the LAN Manager also has a local operator interface, allowing active management of the LAN. Because communications with the LAN Manager use the logical link connections, an implementation of the LAN Manager common to both the Token-Ring Network and the PC Network is possible. The Alert interface of the NetView/PC Program forwards selected error information, compiled by the LAN Manager, to the NetView program. The Alert is displayed at the NetView program console along with probable cause(s) and recommended actions. The NetView/PC program also provides other interfaces to the service point application, such as the Service Point Command Facility (SPCF), which receives commands from the local point for response by the service point application. The SPCF will allow remote LANs and other systems to be incorporated into an automated

operations strategy, whereby human responsiveness is augmented and enhanced with programmed control.¹⁸

REFERENCES

- ¹C. Zarley, "Software Helps Managers Track LANs," *PC Week* (Connectivity Section), January 1989, page C/11.
- ²IDC, "Quantitative Analysis of Local Area Network Markets," *Report IDC # 3173*, August 1987.
- ³M. Pyykkonen, "Local Area Network Industry Trends," *Telecommunications*, October 1988, pages 21-28.
- ⁴P. Schnaidt, "Smorgasboard: 20 Ways to Feed a Hungry LAN Manager," *LAN Magazine*, June 1988, pages 84-95.
- ⁵ISO/IEC JTC1/SC21, *Information Retrieval, Transfer, and Management for OSI: OSIRM Part 4—OSI Management Framework*. Revision of DIS 7498-4, October 1988.
- ⁶M. Hurwicz, "NetView Now," *LAN Magazine*, April 1988, pages 76-79.
- ⁷B. Enyart, "Network Managers' Wish Lists Keep LAN Vendors on their Toes," *PC Week* (Connectivity Section), March 13, 1989, pages C/33-34.
- ⁸J. Schwartz, "Fixing a LAN," *LAN Magazine*, March 1988, pages 70-75.
- ⁹J. Diehl, "Network Tune-up," *LAN Magazine*, May 1988, pages 120-123.
- ¹⁰M. Mohanty, "Troubleshooting Tools," *LAN Magazine*, March 1988, pages 64-69.
- ¹¹E. Ericson, L. Ericson, D. Minoli, *Expert Systems Applications to Integrated Network Management*, Artech House, 1989.
- ¹²M. Hurwicz, "The Sniffer Threat," *LAN Magazine*, April 1988, pages 90-93.
- ¹³E. Tittel, "The LANalyzer," *LAN Magazine*, December 1988, pages 89-93.
- ¹⁴T. Ooka, "Accurate Loss Measurement for Optical LANs," *Telecommunications*, July 1988, pages 48-56.
- ¹⁵R. Watson, "Fortifying a LAN," *LAN Magazine*, October 1988, pages 51-56.
- ¹⁶P. Schnaidt, "The Arsenal: 36 ways to Arm a LAN Manager for Network Battle," *LAN Magazine*, December 1988, pages 69-84.
- ¹⁷M. Willett, R.D. Martin, "LAN Management in an IBM Framework," *IEEE Network*, March 1988, Vol. 2, No. 2, pages 6-12. □

Managing Local Area Networks: Accounting, Performance, and Security Management

This report will help you to:

- Use accounting management to establish charges and identify costs for the use of a LAN.
 - Effectively evaluate the performance of a LAN.
 - Identify potential security breaches on the LAN and take steps to eliminate them.
-
-

With the proliferation of local area networks (LANs), network managers are looking for reliable tools to undertake accounting, performance, and security management.

This report examines techniques, issues, and tools applicable to accounting, performance, and security management.

ACCOUNTING MANAGEMENT

Accounting management enables network managers to establish fees for the use of communications resources, and to identify the cost of using those resources. Generating a report detailing each user's access activities is an important tool for the LAN manager for at least three reasons.

- It can be the basis for chargeback activities.
- It provides a hard copy list of all activities for that reporting period; this can also be employed for security monitoring. If subtle security problems arise later (e.g., discovery of sabotage of data) this type of report may aid in determining what happened and who did it.

This Datapro report was developed by Dan Minoli, an adjunct professor at New York University's Information Technology Institute. Mr. Minoli is also a full-time data communications researcher and strategic planner.

- It can be a basis for LAN usage statistics, an aid in planning for network expansion.

Such statistical reporting features, however, are lacking in all but a few of the currently available network operating systems.

Software Meters

PC software vendors are particularly concerned about product licensing; they want to keep a tight control over the illegal copying of their software. Obviously, a vendor does not want to sell one copy of a product to a large company, only to have hundreds of employees copy it. On the other hand, it would be quite inefficient for the LAN manager to buy as many copies of the software as there are employees. At any time, it is unlikely that thousands, or even hundreds of people need a given application.

Vendors have devised software meters as a way of deterring illegal copying. Certain PC applications now come with built-in meters. A meter works on the same principle as a lending library. When a user starts an application, he/she checks it out of the license library and returns it when finished. If all copies are lent out, the meter returns a temporary denial.

The LAN manager may consider using meters to monitor and control software usage, as well as to provide users with enough copies without purchasing

Managing Local Area Networks: Accounting, Performance, and Security Management

volumes of software. The LAN manager can use the meter to keep usage statistics; this can assist the "traffic engineering" of the software library. Meters range in price from \$200 to \$2,000.

Audit Trail Management Tools

Audit trail systems are a key component of security management, although the function can be considered part of accounting management. Audit trail systems provide information on user activity on the LAN. Additionally, these systems can assist in billing management by providing the data needed to charge back usage.

For an audit trail system to be effective, it must provide the LAN manager with streamlined and useful information (rather than mountains of raw data). The manager may wish to audit only certain users; operations on files with certain extensions or in certain subdirectories; only certain types of operations; or certain servers.² All file and directory creations, deletions, and renames may need to be reported. A system error log report, listing all system error messages to alert the LAN manager of potential problems, may be advantageous. A sophisticated audit tool must allow for this management flexibility. These tools range in price from \$300 to \$700.

PERFORMANCE MANAGEMENT

Performance management enables the user to evaluate the behavior and effectiveness of resources and related communications activities. It is concerned with the use of network resources and their ability to meet user service level objectives. Proactive management can optimize the network's performance. In the past, when smaller LANs (typically fewer than 25 users) were common, performance optimization could be accomplished more easily. Today, at a minimum, a dedicated network manager is needed to manage either large LANs or the connectivity of multiple LANs into a seamless network for as many as several hundred users. LAN performance management also requires a distinction between the backbone LAN facility and smaller sub-LANs in order to monitor and maximize total network use. A company may have either a small number of large LANs or many small LANs connected via a backbone LAN. In either case, users require functionality across multiple bridged LANs to gain access to information that is not directly on their own LAN.³

The first step in improving a LAN's performance is knowing *what* needs to be improved. End users can be a source of information about what needs im-

provement; however, analytical tools are always better. It is not difficult to measure the performance of distinct standalone LAN components objectively. For example, the speed of a file server can be estimated by its CPU type and disk speed. Also, the speed of a LAN transmission medium can be measured in raw megabits per second, which are intrinsic with the underlying technology. User workstations can be measured by how fast a user's screen is written to, or how fast database records are processed.

When one combines all these components together, however, the performance of the LAN is by no means trivial to compute. Simultaneously improving the performance of the file server, the LAN hardware (the network interface cards and cables), and the workstations can assure improvements in the entire LAN system. The situation is more complicated when the user must change only some components of the LAN. For example, by changing settings on the file server, the LAN may seem slower; by improving workstation performance, the entire LAN may seem faster. A reliable methodology is required by the network manager to study these performance issues. Performance management allows the manager to:

- establish a benchmark to measure the current network performance against future measurements;
- provide means to recognize mismatched network or applications software; and
- establish analytical measures to compare possible network performance improvement strategies.

The LAN manager can improve the relative speed at which files are saved, the amount of time spent changing menus, and the amount of time it takes to load programs or data. Several performance measurement programs are now available to determine how well the LAN can move data between file servers and workstations.

Some systems send out records to a file on the file server. Several workstations on the LAN can run the measurement program simultaneously, and the record size can be adjusted to show when the most data is being transferred. Depending on the network software, either the workstation packet buffer size, the number of buffers, or the file server cache block size or packet buffers can be adjusted. By running such programs several times the manager should be able to reach "best case" settings, although this can be a rather long process and requires extensive record-keeping. Other systems are easier and more complete. These tools are a more rigorous approach to measuring LAN performance. By simulating several user applications, such as word processing, spreadsheet, and

Managing Local Area Networks: Accounting, Performance, and Security Management

database entry, and running automatically from several workstations, these programs measure LAN performance in terms of the length of time it takes to do a specific set of tasks. Since these programs keep their own statistics, they are much easier to use. Each workstation's execution time is reported to a text file, and the results are also provided in a graph.⁴

Disk Usage Monitoring

Utilities are needed to monitor disk usage for those operating systems not providing detailed information (such as NetWare). This helps the LAN manager assess whether a user is monopolizing the server and facilitates configuration management in terms of sizing out file server needs. Additionally, the manager will be able to forecast the need for new facilities at future times. Disk usage statistics include the number of files that are in a given directory or volume, the owner of the file, the size of the file, and access chronology. Exception reporting for users colonizing more than a specified threshold of space is available with some products. Also, some products allow the end user to check disk usage. Other products provide partial NetWare security reports by listing users, access privileges, and group membership. These tools generally range in price from \$100 to \$200.

Traffic Monitoring Tools

To undertake performance management, a network manager needs a complete matrix of the LAN traffic patterns. With this information the manager can subsequently look at the LAN configuration to determine, for example, if a server is being used too heavily, or if the network should be partitioned using bridge technology. A protocol analyzer provides detailed information about packets and related Protocol Data Units (PDU) headers which populate the LAN transmission medium, and can compile traffic matrices. Protocol analyzers, however, are complex and need a certain sophistication on the part of the LAN manager to use effectively; and they are relatively expensive.

Software-based traffic monitors will collect some of the needed traffic statistics. In addition to traffic collection, some of these systems can send probes to diagnose nodal problems. Some systems log errors, and issue an alarm when a user-selected threshold is exceeded. A typical performance tool monitors traffic and records how much data is sent to and received from every network node, documenting the packet size, frequency, and type (data or system packet). For system packets, these monitors typically distinguish between commands and internal operations mes-

sages. The data should be collected into sequential ASCII files or a spreadsheet file. Their cost ranges from \$200 to \$8,000.

Issues Affecting Performance

As indicated, file server performance is critical. The file server is shared by all network users, supplying files and applications. As such, it is a major factor in LAN throughput. File server hardware can range from a PC with a small hard disk for minimal use, to a high performance, high-cost, 80386-based system with large and fast external disk drives.

Commonly available file servers (such as those from 3Com and Banyan) are either optimized IBM PC AT compatibles or proprietary microcomputers. In both cases they consist of two basic elements: the CPU, (together with RAM and operating system software) and the disk drive subsystem. If either of these two elements performs poorly, it will impact the overall performance of the file server.

Many LAN network operating systems for printer sharing or program loading may perform relatively well with an 8088-based file server. If the server supports manipulation of large amounts of data, as in a multiuser accounting or database system, then a dedicated file server based on a high-speed processor offers a better solution. Advances in disk technologies have led to 150M-byte drives and controllers capable of delivering a throughput of as much as 500K bytes per second. While these components are not inexpensive, the cost is amortized over many users.

The disk server CPU and the support hardware (RAM) clearly affect performance. If the file server runs software for the 8088 CPU, then RAM cannot normally exceed 640K bytes. Some network operating systems can take advantage of expanded RAM cards to work above 640K bytes. File servers with 80286 or 80386 processors have more RAM available, which will improve file server performance. These systems make more effective use of RAM with file caching and directory hashing, which allow data normally stored on disk to be loaded into RAM. In addition, the file server can access the RAM data in nanoseconds. This allows faster execution, faster table interrogation, and expeditious data movement.⁴ The disk drive subsystem is also important in its own right. Many subsystems (e.g., Novell) use a proprietary disk file structure to achieve higher performance than DOS. This disk format can improve throughput as much as 50 percent. Also, the larger the disk capacity, the quicker the disk drive can move data to the file server. The IBM PC AT disk subsystem is relatively slow, primarily due to the method used by the disk

Managing Local Area Networks: Accounting, Performance, and Security Management

controller to talk to the CPU. With advanced SCSI controllers and high performance drives, data can be delivered much more rapidly. A PC AT with a SCSI disk controller may now be able to move 250K bytes per second or more. Many vendors now use the SCSI disk drive interface.

One bottleneck in the process, however, occurs when the data comes from the disk too fast for the CPU to keep up. A solution is to use a disk co-processor that can hold onto the remaining data until the CPU catches up. Another bottleneck occurs when the data or program files on the drive become fragmented. Most new files are laid out on the disk in a contiguous order, but as time goes by, portions of the file may be rewritten to different sections of the disk. This results in the disk drive taking longer to collect all parts of a file. The easiest way to solve this problem is to do a periodic, complete backup; reinitialize the disk drive; and restore the files from the backup tape.⁴

Interface cards can also affect performance. Memory management is crucial to speed and performance. Factors such as DMA versus shared memory, and on-board processors and buffers can mean large differences in two cards' actual throughput on the network. The performance difference between Ethernet cards can be as high as 50 percent.⁴

Another element affecting performance is the network workstation. The performance of a workstation has more impact on both the perceived system performance and the actual system performance than any other component. For example, a high-performance file server on a 10M bps LAN will show the inefficiency of an IBM PC workstation with limited RAM—the workstation is now the bottleneck since it cannot accept or display data as fast as the file server and the network hardware can supply it. At times, it is cheaper and more practical to upgrade to the workstation, rather than the LAN itself. Adding more RAM or a co-processor could improve grade of service without a single change to the network. The protocol software can also affect workstation performance. A full seven-layer OSI (Open Systems Interconnection) stack could require considerable resources to run. Even at the Network layer, packet sizes, transfer buffers, and other workstation network software settings can have a major effect on the network performance.

SECURITY MANAGEMENT

Security on LANs has become an important issue recently, primarily due to the publicity given to network viruses. LAN security issues involve tapping, radiation leakage, user authentication, file and pro-

gram security, audit trails, encryption, and physical security. While security management in mainframe computer systems is well-developed and mature, LAN security is still in its infancy, according to industry experts.

Security management supports the application of security policies. It controls access to both the network and the network management systems; it may also protect information transacted by the LAN from disclosure or modification.

No product (or product family) on the market today provides the total solution for maximum LAN security. However, there are some relatively good LAN software and hardware products available that will meet the security requirements of most installations. Secure LANs are configured as a combination of products from several vendors.

Choosing a LAN software or hardware configuration that will support the needed security requires an understanding of LAN architectures, the safeguards provided by the products, and the specific security requirements of the target environment. With the network operating systems now available, centralized data storage is achieved via common access to the file server. In this environment, critical files may reside on a central device that is accessible by all workstations within any workgroup. Unlike mainframe database files, however, there are many considerations that must be addressed by the LAN manager to insure proper security. Control of access to LAN data poses several security problems that are not found in mainframe installations.¹

- PC LAN users are sometimes more sophisticated than "dumb" terminal users. Because the PC operator must acquire some knowledge of DOS and its commands, an understanding of internal security structures is more common.
- In the mainframe environment, only computer operators have access to tapes and hard disks; in a LAN environment every PC (except for diskless workstations) stores data. Protecting this data from theft or destruction becomes a more difficult task.
- Utilities are readily available to do bypass copy protection, expose disk substructures, and perform sophisticated file/disk copying. Use of these utilities exposes all data on the local workstation or LAN file server to security risks.

Evaluation of the company's (or even workgroup's) security requirements is important in making the decisions pertaining to LAN configurations—low security LANs are generally less expensive and allow the

Managing Local Area Networks: Accounting, Performance, and Security Management

choice of a wider selection of software and hardware, while high-security requirements may force the selection list to only a few options. Additional hardware and more expensive software is generally required for the more secure LAN installations. Implementation of a very secure LAN is considerably more costly than a single workgroup LAN. LAN security issues fall into three major areas.¹

Physical Access. Security in any data processing environment starts with controlling access to the equipment. Although intrinsically distributed in topology, LAN security requires installing the file servers and printers in secured access rooms. Access to the LAN's cabling system is also a concern because of the potential to "tap" into the network, insert new nodes, or monitor network data traffic. Access control to the PC workstation itself must be considered. Even without the network or server available, the local hard disk of the workstation can pose a security risk for loss of data.

Logical Access. Physical access techniques are designed to keep unauthorized users off the network. Logical access techniques are designed to keep authorized network users away from unauthorized files. Access to the data is the responsibility of the network operating system. It is via the NOS that the logical control for information access is carried out. Password access to servers, I/O rights to directory or file structures, and user accounting features represent typical support features provided by a network operating system. The level of security required by any site will dictate the LAN manager's final choice of LAN software.

Administrative Control. An important but often neglected aspect of LAN security is the role of the LAN manager. It is this individual who is responsible for physical and logical access control, in addition to undertaking fault recovery procedures, performing backup, and monitoring for potential security infractions.

Physical Access

Physical access security for LAN installations, in turn, affects the three major components of a LAN: the workstations, the servers, and the cabling. Workstations may represent the highest security risk in a LAN—unchecked access to workstations and their local storage provides an avenue for theft of sensitive information. An authorized user may download information from the server to the workstation; once stored on the local disk, no security processes are in place to prevent an unauthorized user from obtaining its contents except by stringent physical controls.

This involves installation of one or more add-on items of software and hardware. For example, unauthorized use of a workstation may be prevented by attaching keyboard lock devices, either software- or hardware-based. With such tools, only the authorized user of that PC can activate the keyboard. Also, adding physical restraint equipment to the PC itself will prevent it from being removed from its work area. Such restraints were common for \$1,000 typewriters, but are not as common for \$5,000 PCs.

Diskless PCs are another means of physically securing a LAN. Diskless PCs use the standard PC AT/XT bus with an important difference—there are no diskette or hard disk drives. Each user stores data on the server's hard disk. Diskless PCs prevent users from stealing corporate information or software. (Other institutions, particularly the government, also eliminate printers.) By eliminating the disk drive, diskless PCs make it difficult to introduce viruses onto the network.

Because of the decreased physical size of most file servers, this equipment is often subject to theft. The file server should be placed in the most secure location possible; this could be an MIS computer room or a special room designed for secure equipment. Even if physical access is strictly controlled to prevent theft of the server itself, loss of information could occur through misuse of the server console. The storing, archiving, and vaulting of the system's backup tapes must also be considered carefully. On-site storage should be under the same security restrictions as the file server, but (if possible) in a separate location for disaster recovery purposes. An off-site storage facility must also meet the same security requirements as the central site. Other areas of concern include communications gateways, which can be accessed through dial-up. Security problems, including virus infections, often arise because of the poor control over these calls. A separate port controller should be considered where security is critical.

Cable is one of the first and easiest places for a LAN security infraction to occur. Copper-based systems can be tapped easily—to tap into a twisted-pair LAN one does not even require direct contact with the cable. An intruder can use an electromagnetic pick-up antenna that costs less than \$20. In spite of what many have been led to believe, optical fiber is also tappable; it just takes more sophisticated equipment and more skill. Fiber is only slightly more secure—a perpetrator can use a razor blade-like knife to peel back the fiber's cladding. The cut is deep enough to penetrate the cladding, but not the fiber itself. Once the cladding is compromised, the intruder can tap the fiber cable. Also, to extend a fiber link there must be a way to connect to the cable. These connections are

Managing Local Area Networks: Accounting, Performance, and Security Management

fiber systems' potential weak point. Devices to tap a fiber connector cost around \$275.⁵ The physical cabling is at risk if unauthorized personnel are allowed to tap it or to attach a network monitor and protocol analyzer. It is possible with these devices not only to determine user passwords by observing packet traffic on the cable system, but to actually capture sensitive information.

Electromagnetic signal leakage outside the building is another vulnerability for LAN security. Coaxial and twisted-pair cabling and the devices which connect cabling (connectors, amplifiers, tap boxes) will leak a certain amount of signal. Depending on the quality of the antennae used, these signals can be decoded from a quarter-of-a-mile to several miles away. Fiber cable is more secure in this respect.

Logical Access Security Issues

Logical access control is provided principally by the LAN's network operating system. Prevention is the first line of defense. An emphasis on prevention will save an organization time and money in the long run.

The security management task is the maintenance of the "access environment," which includes all aspects of access control, including the methods of control, monitoring the effectiveness of the control, and reporting the saved audit trails for later analysis.

A common way to attack LAN security is through a PC on the LAN. LAN security schemes must include ways of controlling access to the network. There are two basic forms of access control: user authentication (usually in the form of password, but also with biometrics) and file and program security.

User Authentication. A user must own a legitimate password to gain access to the system. Theoretically, an unauthorized user will not have a valid password, and thus be excluded from the system. Unfortunately, passwords are notoriously weak. Users tend to employ simple-to-remember household names. They are written down, sometimes right next to the terminal; they may be shared with other users; they remain unchanged for long periods of time; and not enough companies have rigid controls to revoke all passwords from all systems when people are terminated or resign. Some systems make users' passwords expire every 30 days. With Novell's NetWare Version 2.1, the LAN manager can set the password expiration period to any length.⁵ Password length and randomness are also critical. There are a number of password generator products on the market.

Onetime password systems are among the more reliable password mechanisms. This approach employs convenient user-owned devices that generate a one-time password for the user to enter into the system, and a LAN-based software counterpart. Both the remote device and the LAN employ the same algorithm to generate the next legitimate password. Access depends on the user's possession of the device. Naturally, the user must safeguard the portable device, as one would safeguard a credit card. In addition to the security risk, losing one of these devices can be expensive—the portable generators cost around \$200 each.

Biometrics is one of the best ways for securing a computer room. Biometric access control is purported to be the best way to verify that the user is really who he/she claims. Biometrics use unique body characteristics for identifying an individual; these cannot be stolen or forgotten. Biometric controls work from the actual physical presence of the user. Devices that read a user's fingerprints or thumbprints can be used. Another product uses infrared light to measure retinal patterns. Biometrics is used mostly in government installations. These devices are expensive, however, and they are still not perfect.

In particular, voice-prints are highly unreliable means of identifying people, because the signal processing techniques used are very crude; some of these systems that can be trained by a particular user do not even recognize that user's voice when affected by a cold. The fingerprint reader records users fingerprints and stores them on file.⁶

At this time, the most reliable method is a combination approach: a password generator or biometric scheme used in conjunction with a memorized password. Many LAN managers, however, now rely only on simple passwords.

File and Program Security. Once a user has been authenticated and has gained entry into the system, security concerns turn to what the user can access. This is called access control in the evolving international security standards. The effectiveness of any access control scheme depends on how granular a control the LAN manager has over resources and objects. In a LAN environment, resources include servers, disk volumes, directories, and files. A user is assigned access rights to one of these objects and that determines the operations that can be performed. A minimum of READ or WRITE access control is given to the user, while some operating systems extend control to UPDATE, ADD, and DELETE. In terms of implementation, access control is provided by a combination of the features of the network operating system and the specific security utility. While

Managing Local Area Networks: Accounting, Performance, and Security Management

the NOS will allow access control down to the directory or file level, the security utility must provide control at the record and even field level.¹

The manager needs to delimit and limit the data a user can look at, and even control the applications and utilities that can be accessed. Most network operating systems offer a form of file security, where the LAN manager can determine what volumes, directories, and even individual files a given user will have access to. The LAN manager can also set the type of operation, determining whether the user can only read a file or whether he/she has full read/write access. Access must also be controlled to programs. One way to prevent unauthorized users from running given programs is to put them in a directory with a file security system that will not allow these users to read them. Another approach is to use an application control system.⁵

Access monitoring allows the LAN manager to track security problems in real time. Most network operating system software lacks this capability, which is common in mainframe systems. The capability to detect a security infraction as soon as it happens is often critical to determining the problem and quickly resolving it. With this information, the manager can temporarily suspend the user account in question, review security detail information, and make a determination as to the source of the problem.

Administrative Issues

Administrative procedures and responsibilities are critical to security. The LAN manager must generate the security policies, perform backup/restore procedures, and implement the specific access architecture to support the desired level of security. This educates all users as to the scope of the security policies. For example, the simple practice of logging out from the network each time a user leaves his/her desk is a key component for maintaining security. Periodically, the LAN manager should run diagnostic utilities to insure that disk information is intact. Also, a review of the actual data itself should be performed, in order to detect possible sabotage or corruption of the information.¹

Maintaining the data integrity of any data processing system requires procedures that provide insurance against disk failures. Usually, backup security involves periodically generating a backup copy of the information. Magnetic tape is typically used for backup because it is inexpensive and can be stored economically. Novell's SFT (System Fault Tolerant) NetWare is the first step in addressing backup and disaster recovery; this version of NetWare can dupli-

cate directory structures, do read-after-write verification, and do on-the-fly disk sector error recovery. However, if the server crashes, the user still needs to reconstruct the data from a separate backup (possibly from tape). The user must manually reenter some bindery information.

The removable nature of tape reels and cartridges represents another security risk. When tape backup is performed, the format of the tape data is in one of two formats: image or file-based. The image tape format is an exact disk image; restore operations are direct, consisting of restoration of the complete disk. File-based tape methods sequentially store entire files in directory order on the tape. With a file-based tape, a certain amount of search is needed to restore files from tape, exposing all files. In considering security aspects of backup/restore utilities, the image format is both faster in execution and more secure because of the limited operations support for these formats. Many backup software packages require that image backup be restored to the same drive they were backed up from.

Fault recovery tools automate rebuilding servers. This is done by gathering critical file server information on hard disk or diskettes, including the bindery, login scripts, directory rights, system auto exec file, and printer definitions.² Products typically range from \$200 to \$1,500.

Protocol Analyzer Issues for Security

A protocol analyzer in the hands of the wrong person can be a security threat. If an infiltrator can get access to a LAN port or is able to tap the cable, the analyzer can reveal useful penetration information. An analyzer can capture the entire dialogue taking place over the LAN, and can display passwords in an easily readable form. Appropriating passwords is easy with analyzers, but passwords may not always be useful in a properly designed LAN. It is possible, for example, to restrict the station(s) a user can log in from; thus, although the infiltrator may have the manager's password, he/she cannot log in as the supervisor without using the actual manager's terminal. In addition, audit trail utilities can report logins and logouts, with special attention paid to the manager's ID. This is the reason why reliable security measures that go beyond basic password protection are needed.

While protocol analyzers can present problems for LAN security, they can in turn be used to monitor the network for infractions. One simple technique involves looking for stations that are not supposed to be on the network. The manager can set the display to depict unknown stations. This is done by declaring an

Managing Local Area Networks: Accounting, Performance, and Security Management

easily readable name for each LAN station. If a program claims to lock certain files, for example, the analyzer can be used to test that claim. With some analyzers, the network manager can write programs in C for specialized functions, such as monitoring compliance with security procedures. For example, such a program might look through the data to find stations that are logged on to a file server but show no activity for long periods of time. This may indicate a station where the user has walked away without logging off, which is a violation of security policies in most institutions.⁷

REFERENCES

- ¹R. Watson, "Fortifying a LAN," LAN Magazine, October 1988, pages 51-56.
- ²P. Schnaidt, "The Arsenal: 36 Ways to Arm a LAN Manager for Network Battle," LAN Magazine, December 1988, pages 69-84.
- ³M. Pyykkonen, "Local Area Network Industry Trends," Telecommunications, October 1988, pages 21-28.
- ⁴J. Diehl, "Network Tune-up," LAN Magazine, May 1988, pages 120-123.
- ⁵M. Mohanty, "Defending a LAN," LAN Magazine, April 1988, pages 84-88.
- ⁶D. Greenfield, "Sensible Paranoia," LAN Magazine, April 1989, pages 84-88.
- ⁷M. Hurwicz, "The Sniffer Threat," LAN Magazine, April 1988, pages 90-94. □

Managing X.25 Packet Switched Networks

This report will help you to:

- Evaluate alternative packet network architectures and the network management issues related to each alternative.
- Handle the added layer of network management complexity created when X.25 technology supplements conventional network architectures.
- Learn who the suppliers of X.25 Network Management products and services are.
- Choose between public versus private X.25 network alternatives.

The use of packet switched networks for data communications has grown enormously during the past decade. These networks are popular because they offer a low-cost way to access one or many host computers from remote locations in a very reliable fashion. The types of networks offering *public* packet-switching services are called Value-Added Networks (VANs).

While many users have utilized packet switching over VANs, others have installed private packet-switching networks to provide internal VAN-like services. This is analogous to using private lines in place of dial-up public access or to using a PBX instead of Centrex service for voice communications.

In 1976, the CCITT introduced standards for accessing public data networks via X.25 lines. The CCITT subsequently revised these standards in 1980, 1984, and 1988.

This report was developed exclusively for Datapro by James B. Wetterau, P.E. Mr. Wetterau is the principal and founder of Networking Solutions, a New York City-based network consultancy. With over 25 years' experience in the communications field, Mr. Wetterau conducts frequent seminars nationwide on various data communications topics. Networking Solutions advises clients in network planning, design, implementation, and management in X.25, T1, SNA, and LAN technologies.

Due in part to the CCITT's initial efforts and refinements, the X.25 link is now the most popular way to link host computers to packet-switching networks. The X.25 line has become so closely associated with packet-switching networks that the term "X.25 Networks" is incorrectly used to describe packet switched networks which use X.25 links to connect to user computers. X.25 is more precisely defined as the interface between data termination equipment (DTE) and data communications equipment (DCE) for packet mode terminals connected to public data networks by private line.

Thus, when we discuss the management of X.25 Packet Networks, we are in fact discussing the management of packet switched networks which make use of

Index to This Report	Page
X.25 Network Architecture	102
X.25 Network Management Issues	103
Managing a Value-Added Network (VAN)	105
Managing a Private X.25 Network	106
X.25 Network Management Products	109

Managing X.25 Packet Switched Networks

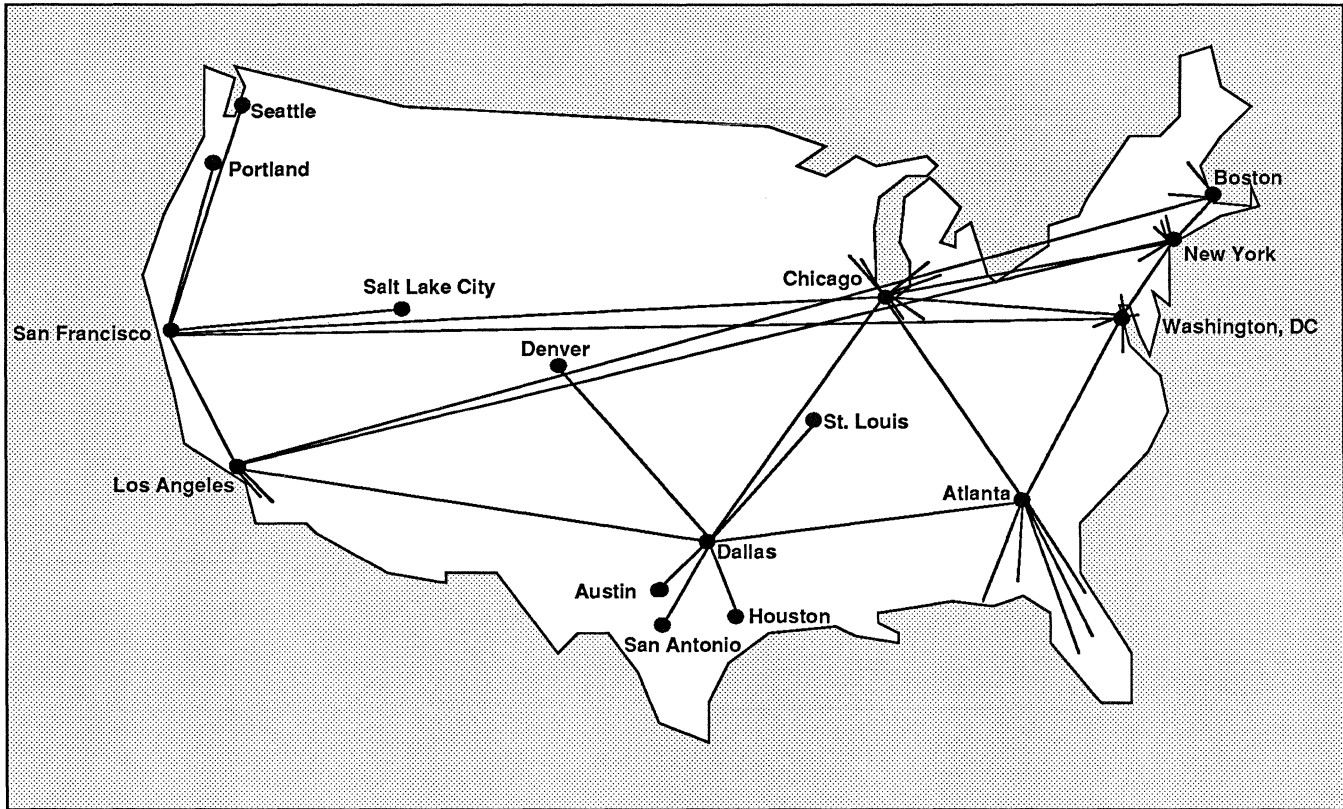


Figure 1. The Telenet backbone network illustrates the redundancy and fail-safe features of X.25 packet switched networks. All circuits shown are 56K bps.

X.25 links to connect user devices to the packet networks. This report briefly reviews the history and technology of these networks before discussing relevant network management issues.

Packet switched networks evolved from a combination of three different technologies: the ARPAnet system, timesharing services, and the CCITT's standards efforts.

During the 1970s, the U.S. Defense Department embarked on a series of projects designed to achieve more effective communications between incompatible computers. To help accomplish this, the Defense Advanced Research Projects Agency (DARPA) set up ARPAnet—a global network that uses File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), TELNET, and other protocols to network multivendor equipment. DARPA's secondary goal was to achieve reliability with message rerouting to provide more robust networks in the event of war.

At the same time, timesharing services were becoming popular among user organizations that were too small to own dedicated computer facilities. As timesharing services grew, users wanted more cost-effective ways to reach the timeshare hosts. Dial-up long-distance calls

grew too expensive, and dedicated lines were usually too wasteful for single-terminal access.

In 1976, the Consultative Committee on International Telephony and Telegraphy (CCITT) proposed the X.25 standard in an attempt to define a more reliable method of data communications in the public (PTT) telephone network.

During the 1970s, all three of these trends came together. Based on the ARPA research on packet routing, the first VAN (Telenet) was established. Subsequently, the timesharing vendor Tymshare, Inc. created its Tymnet packet switched service to allow more efficient access to timesharing services. Both VANs adopted the X.25 standard as a more efficient way to connect computers and users' terminals to their networks.

Packet network growth exploded during the 1980s, with IBM (Information Network) and AT&T (ACCUNET) offering packet-switching services. Figure 1, which shows the Telenet backbone cities, illustrates the nationwide coverage of VANs.

Managing X.25 Packet Switched Networks




Figure 2. A packet switched network uses X.25 lines for dedicated access and dial-up lines for short-term access.

X.25 NETWORK ARCHITECTURE

X.25 packet switched networks exhibit the general structure shown in Figure 2. The network employs switching nodes connected by private lines. The user host computer is connected to the network by an X.25 link, which supports multiple virtual circuits and packets between the host and the network over the same physical connection. Terminal connections are supported either by dial-up into a local network node or by private X.25 lines, usually with a packet assembler-disassembler, (PAD). Messages are broken into packets that are transported through the network. The path for each message is not fixed, and packets for the same message can take differing paths, depending on congestion. Packets are reassembled at the destination node. The actual network topology depends on total network traffic and is transparent to the user. The only requirement is that each node has more than one link so that traffic can be rerouted in the event of failure. Packet-switching networks are designed to allow various kinds of hosts and terminals to communicate with each other. Thus, most nodes support the following types of protocols:

- Asynchronous
- Bisynchronous
- Synchronous Data Link Control (SDLC)
- Specialized file transfer
- Electronic mail

These various incompatible protocols do not become compatible over the packet network, however. The packet network provides a transparent highway over which differing protocols can share transport capability.

To properly manage an X.25 network, however, a network manager must understand the X.25 protocol. The X.25 protocol specifies the kinds of packets required to transfer information from the user of the packet network to the final destination.

X.25 PACKET STRUCTURE

X.25 has three components or layers:

- **Layer 1**—the physical interface, X.21 and X.21 bis, which are similar to RS-232.
- **Layer 2**—the link access protocol, which is Higher Data Link Control, Link Access Protocol Balanced (HDLC Lap B). The structure of HDLC is similar to SDLC.
- **Layer 3**—the X.25 packet layer, which contains the rules for successful data transfer from user DTEs across public data networks.

The X.25 packet structure is shown in Figure 3.

The X.25 packets are made up of octets (bytes) that control the data flow across the network. Four bits of the first byte are the general format ID. This determines whether it is a modulo 8 or 128 packet network, as well as the network signaling requirements for the Q and D bits. The other 4 bits of byte 1 and all of byte 2 make up the 12 bits of the logical channel address. Twelve bits are required to accommodate the 4,095 virtual circuits allowed on the X.25 link. Byte 3 defines the packet counters that control the message flow. If it is a modulo 128 network, Byte 4 also contains packet counters. The remainder of the packet is user information.

Users may implement this architecture either with VANs to provide packet-switching services or by building their own packet network. The three major VANs in the United States are US Sprint Telenet, McDonnell Douglas Tymnet, and AT&T ACCUNET Packet Service. In addition, the Bell Operating Com-

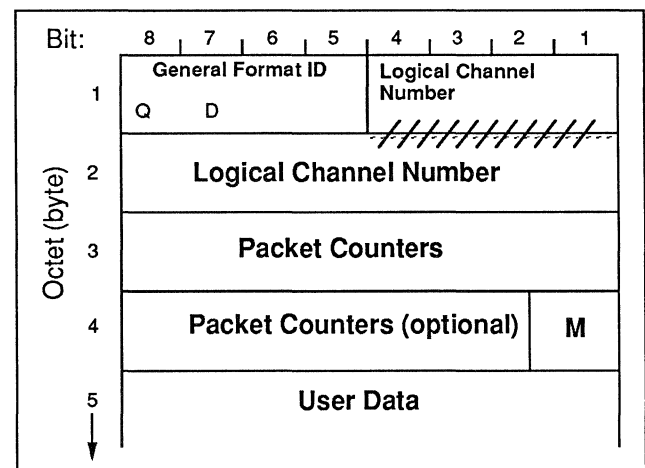


Figure 3. The X.25 packet structure uses three bytes for modulo 8 and four bytes for modulo 128 packets.

Managing X.25 Packet Switched Networks

panies (BOCs) have recently implemented packet networks in their serving areas.

If users build their own private packet network, they must install both network access facilities and packet-switching nodes at the locations to be served by the network. User terminals and hosts would then be connected to the node at their location.

The decision of private versus public network access is an economic one, just as with other types of data network services. Typically, for lower levels of use, public access will suffice. If network usage cost continues to increase, a private network may be cost justified. Both architectures are discussed in the following section.

X.25 NETWORK MANAGEMENT ISSUES

The management of packet-switching networks encompasses many of the same factors as managing any data network. The X.25 parameters, however, add a layer of complexity to the overall network management process. X.25 standards are similar, though not identical, to the first three levels of the OSI model (i.e., levels 1 through 3). Thus, the X.25 network structure attempts to emulate the network management structure of OSI. This structure calls for Specific Management Functional Areas (SMFAs), defined as:

- Configuration Management
- Fault Management
- Performance Management
- Security Management
- Accounting Management

For more information on OSI specifications covering these five SMFAs, see "OSI-Based Network Management," Report NM40-200-101.

X.25 networks require all of the above areas to be accounted for in the management tasks.

CONFIGURATION MANAGEMENT

X.25 circuits require the specification of configuration parameters. These parameters indicate the network address, how many simultaneous virtual circuits are allowed on each line, and what alternate routing will be specified in the event of circuit failure. The network manager specifies the X.25 packet fields for the network's nodes and lines.

FAULT MANAGEMENT

Fault management in X.25 networks is similar to that of any synchronous protocol network. The primary

difference is that an X.25 packet switched network guarantees network delivery of packets. Thus, in the event of failure of hardware or software, packets must be alternate-routed to guarantee delivery.

From the hardware perspective, it is desirable to build in enough redundancy so that any failure does not prevent message delivery. This often involves duplicate access to packet-switching nodes or duplicate circuits. In addition, for private networks, the nodes will be configured with duplicate components and power supplies, as well as alternate routes for circuits.

Software fault management involves detecting missing packets and incorrectly delivered messages. This primarily requires the ability to understand the X.25 protocol, which is similar to understanding SDLC in IBM synchronous networks. To properly understand X.25, one must understand both the HDLC Lap B structure, shown in Figure 4, and the X.25 packet format illustrated in Figure 3.

HDLC Lap B has the same fields as an SDLC frame. A uniform eight-bit flag (01111110) starts and ends the frame. An address field of eight bits routes the packet. The eight-bit control field defines the type of message. The 16-bit FCS provides error checking. The X.25 packet is inside the LAP B frame, in the variable-length information field. Lap B defines the structure of the messages transported across a single link between the user and the network. Lap B messages are of three types, defined in the Control field:

- Information transfer, (i.e., data packets)
- Supervisory messages
- Unnumbered frames

Data packets are transmitted in information frames, while supervisory and unnumbered frames are used for call setup and supervision. Lap "balanced" (B) is used to indicate that a balanced mode of communications is used, where either end of the link may transmit at any time. This is referred to as BA mode for balanced asynchronous (since neither node is in charge). BA mode requires more elaborate control software at both nodes on a link but provides greater efficiency and control.




Figure 4. The HDLC format has the packets in the information fields.

Managing X.25 Packet Switched Networks

Packets inside the Lap B frame perform the functions shown in Table 1, which include:

- **Call setup and clearing.** The configuration is set up with Call Set Up packets before the virtual circuit is established. These packets select the route, as well as the alternate route in the event of failure.
- **Data and clearing packets.** Data packets route the necessary information across the network, while interrupt (clearing) packets are provided for circuit emergencies.
- **Flow control.** Flow control and reset packets tell if the DCE or DTE is "ready" or "not ready" for data transfer. This aids in routing.
- **Reset and Restart.** Restart packets reestablish a call, and registration packets are part of the audit trail.
- **Diagnostic and Registration.** Diagnostic packets aid in troubleshooting and error recovery.

Traffic monitoring permits control of utilization for both circuits and nodes; it also provides the information required for security and billing procedures. From these statistics overall network use can be determined and growth planned. Finally, accounting procedures

include billing based on traffic and location. Thus, traffic statistics are required to adequately provide accounting services.

MANAGING A VALUE-ADDED NETWORK (VAN)

VANs are the primary means by which most users access packet networks today. In this type of configuration (shown in Figure 5), the following management criteria are most critical to ensure successful operation.

Vendor relations are particularly important when VANs exercise such major control of your network. The network control and management issues should be addressed and procedures spelled out in advance of going "live" with the system.

A user's host computer is connected to the VAN via private lines to the nearest VAN network central office. Typically, the VAN procures these lines from the user and manages the lines from the VAN Network Control Center (NCC).

An X.25 software module in the user host (as indicated for Host A) communicates with the packet network and controls the flow of information to and from the VAN. Up to 4,095 simultaneous connections are possible on this X.25 link. To preserve this traffic flow,

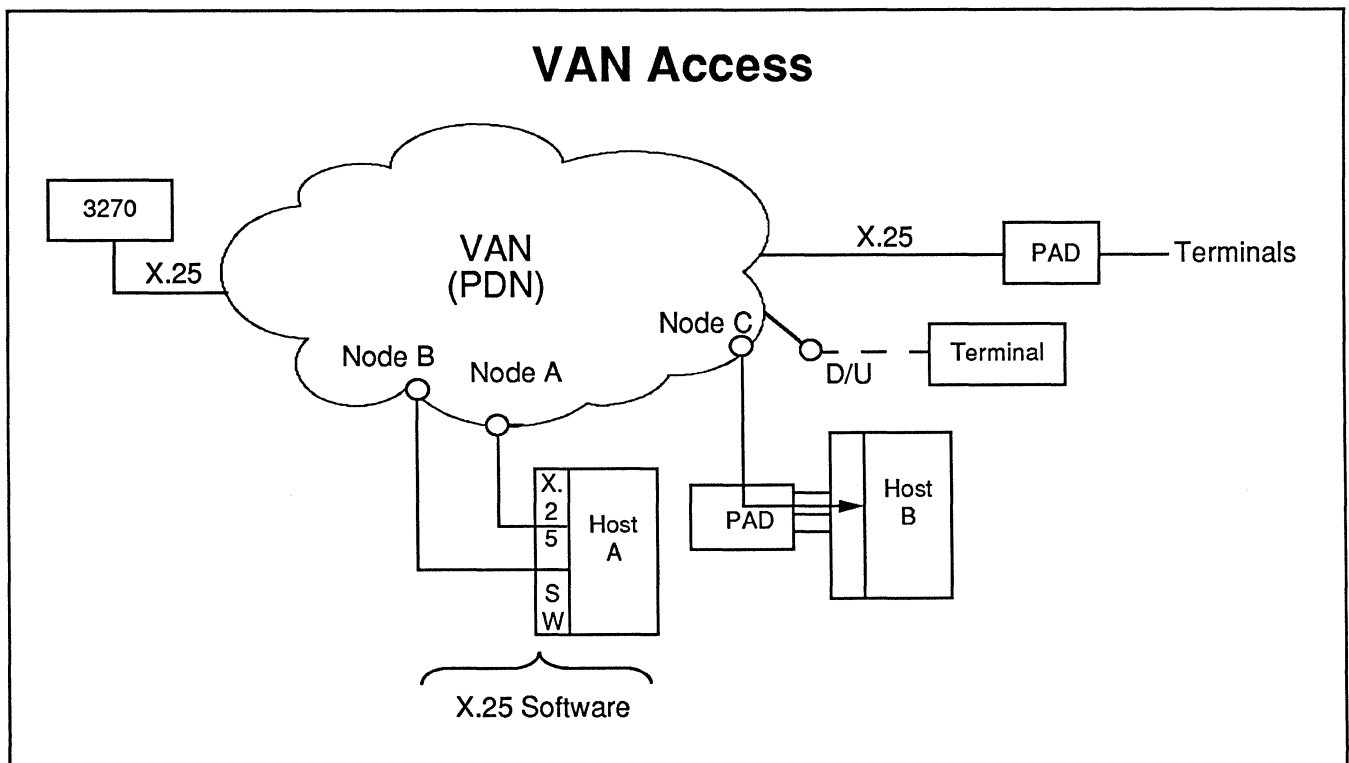


Figure 5. A typical configuration for VAN-provided packet-switching access.

Managing X.25 Packet Switched Networks

DCE to DTE	DTE to DCE
Call Connect and Disconnect	
Incoming Call	Call Request
Call Connect	Call Accept
Clear Indication	Clear Request
DCE Clear Confirm	DTE Clear Confirm
Data	
DCE Data	DTE Data
Interrupt	
DCE Interrupt	DTE Interrupt
DCE Interrupt Confirm	DTE Interrupt Confirm
Flow Control	
DCE Receiver Ready	DTE Receiver Ready
DCE Receiver Not Ready	DTE Receiver Not Ready
	DTE Reject
Reset and Restart	
Reset Indication	Reset Request
DCE Reset Confirm	DTE Reset Confirm
Restart Indication	Restart Request
DCE Restart Confirm	DTE Restart Confirm
Diagnostic and Registration	
Diagnostic	Registration Request
Registration Confirm	

Table 1. X.25 packet types.

either dial backup or a duplicate line is desirable. A duplicate line is definitely preferred if it is critical for the user's network to keep functioning. In the event of line failure, the X.25 software in the host and the routing tables in the VAN node can detect the failure and auto switch all traffic to the other line. If a host does not support X.25, a PAD can connect the network via X.25 and to the host asynchronously (as shown for Host B).

To increase failure protection, the two circuits should be routed to two different VAN central offices (COs). This will protect against node as well as circuit failure. If the user has dual CO connections, the routing tables can automatically send all traffic to the alternate CO in the event of node failure.

FAULT MANAGEMENT

The VAN's NCC performs network monitoring for the user and assumes total control of the user configuration. In the event of a circuit failure, the VAN NCC would report it to the appropriate carrier. Nonetheless, since many VAN COs are unattended and remotely monitored, it is advisable for the user's network man-

agement system (i.e., matrix switch, patch panel, etc.) to also monitor the X.25 circuits. If the modems indicate loss of carrier, the user can detect it as quickly as the VAN NCC staff.

Trouble detection is thus primarily the function of the VAN network management. Users will be able to see any modem or DSU indicators, possibly including carrier loss. If any failures are detected, users should call the VAN trouble number, if the VAN network management staff has not already contacted them.

VAN vendors have a standard mode of reporting troubles to the user. The network manager should provide the vendor with contact names and numbers for 24-hour access. Be aware that if the user node location is unattended at night, failures will not be fixed until the following working day. The network manager should obtain vendor names and phone numbers, including an escalation procedure for unresolved problems. If lines or nodes are down, the user's entire network is inoperable. It is imperative that the vendor response be timely and appropriate.

CONFIGURATION MANAGEMENT

Properly configuring X.25 software parameters is important in packet switching via VANs. The user specifies configuration parameters when loading the X.25 software for initial setup. The VAN support staff should help in this process.

Of primary concern is the number of simultaneous virtual circuits. This quantity should be high enough to avoid any busy-hour call blockage. Since users usually dial in at 300 or 1200 baud, a 9600 bps X.25 line can usually handle 32 simultaneous connections with no difficulty. If utilization exceeds 75 percent, it is time to consider another line.

Users can also monitor the X.25 software installed in their host. The host console will provide statistics about the X.25 line performance, including addresses, traffic, busy lines, and alarms. For IBM networks, the X.25 software is installed in the 37X5 front-end processor. The two primary software packages used for this are Network Packet Switching Interface (NPSI) from IBM or the X.25 package from Compro Associates. Both of these packages provide console support for network monitoring.

PERFORMANCE AND SECURITY MANAGEMENT

VAN vendors also provide operational statistics about the network on a monthly basis. These statistics will indicate the number of users, where they called from, how long they were connected, and how much traffic

Managing X.25 Packet Switched Networks

they introduced into the network. Your bill is based on this information. It is useful to assess this data regularly to determine areas for network improvement.

The network manager should work with the VAN vendor to establish a well-defined set of procedures for those who access your host via the VAN. It is the host system provider's responsibility to inform those users accessing the host about procedures for network access. Typically, the VAN vendor provides a booklet that lists all VAN access telephone numbers nationwide and briefly explains logon procedures. The host system provider must assign logon passwords and IDs to prospective terminal users for security. These are assigned via the X.25 control console. Well-written instructions will save many hours of phone consultation with the host system users.

COST CONTROL

Cost control is a major concern for packet switched networks. When access to your computer is available with a local call, at no perceived cost to the caller, abuse is possible. The VAN software will automati-

cally disconnect users who exceed a predefined time limit with no activity; this is usually a desirable feature.

The best way to control network access is with a chargeback scheme. When long-distance access to your host is provided with a local phone call, it is appropriate to replace part of the savings in calling costs with a chargeback procedure. Charges can be based on time connected or traffic or both, although it is usually a function of connect time. Chargeback schemes help organizations recover some of the VAN costs and, more importantly, those who access the computer are made aware of the associated cost of this improved access.

MANAGING A PRIVATE X.25 NETWORK

The level of management complexity and control required increases significantly when a user implements a private packet switched network instead of using a VAN. All the network control activities that the VAN vendor carried out must now be done by the user who controls the network directly.

Fortunately, the software needed to carry out this degree of control is resident in the hardware nodes.

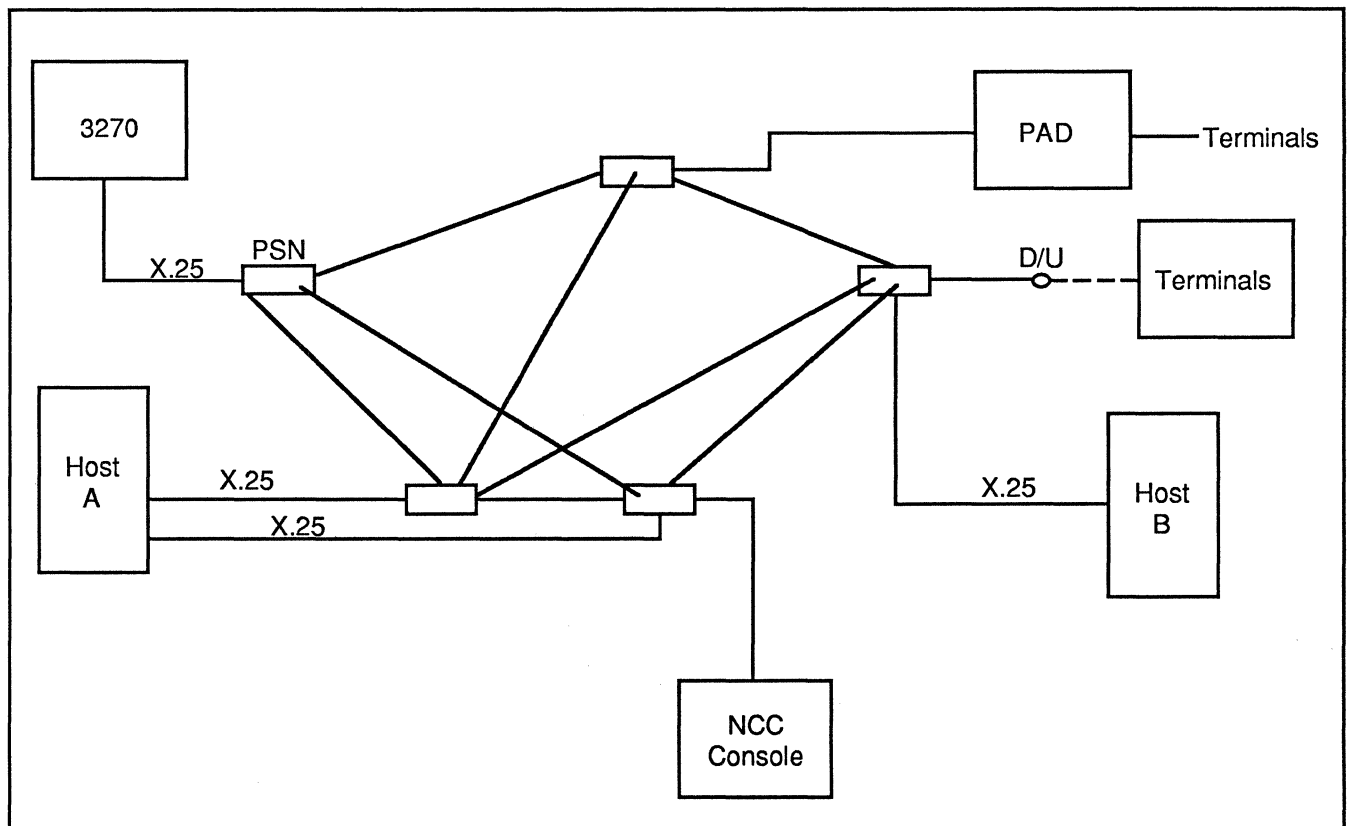


Figure 6. A sample configuration depicting private packet-switching access.

Managing X.25 Packet Switched Networks

Network nodes require routine maintenance, as does any piece of hardware. Routine maintenance is typically performed by the hardware vendor under maintenance contract. The hardware vendor should provide spare components, either ports or entire backup nodes, so that all network addresses and routing instructions can continue executing on the backup hardware during scheduled preventive maintenance periods. Since one of the hallmarks of a packet network is its fail-safe performance, it would be distressing if insufficient redundancy were provided to allow for routine maintenance.

Dual circuits should be provided from each node to ensure reliability in the face of circuit failures (as shown in Figure 6). In selecting circuit quantities and locations, it is desirable to minimize the total number of hops between any node destination points for users. The user's network operations staff is responsible for performing circuit monitoring and follow-up with carriers. Thus, all circuits should be well documented. In addition, selecting circuits from alternate vendors for reliability is desirable, assuming several vendor choices are available.

FAULT MANAGEMENT

Troubleshooting within a private network focuses on the hardware nodes and circuits connecting them. The network manager should establish a Network Control Center (NCC) for monitoring all hardware and software components and circuits. Matrix switching or patching access should also be used for all circuits. Remote monitoring of all nodes with alarms to the NCC should be provided. Besides following normal guidelines for network maintenance for any other type of network, network managers should also use two additional primary tools for troubleshooting:

NCC and System Console. Typically, packet network hardware includes an NCC system and console at which error messages will appear. This is the network manager's most valuable detection device for fault isolation. The packet network NCC collects all types of statistics for network management. The console should be readily accessible to the support staff who must manage the network. In addition, if there are large quantities of circuits, a network console from the modem vendor is a useful tool since a private network requires that users support their own DCE devices as well.

X.25 Data Analyzer. An X.25 data analyzer is indispensable for troubleshooting problems not readily diagnosed by the circuit monitor or packet network console. Data analyzers enable the network manager to

monitor traffic to ascertain correct performance; they also identify packet errors, based on prerecorded correct X.25 information.

Furthermore, the data analyzer can act as a protocol simulator, allowing one-shot testing of devices, to verify errors or correct performance.

TESTING X.25 SOFTWARE

The ability to test the X.25 software is a major part of the NCC function. Diagnostic packets provide information to troubleshoot software problems. Routing tables and alternate paths must be defined to ensure

Optional Contracted User Facilities

- Extended packet sequence (modulo 128)
- Nonstandard default window size
- Nonstandard default packet size
- Flow control negotiation
- Throughput negotiation
- Packet retransmission
- Incoming calls barred
- Outgoing calls barred
- One-way logical channel outgoing
- One-way logical channel incoming
- Closed user group
- Closed user group with outgoing access
- Closed user group with incoming access
- Incoming calls barred within a closed user group
- Outgoing calls barred within a closed user group
- Bilateral closed user group
- Bilateral closed user group with outgoing access
- Reverse charging acceptance
- Fast select acceptance
- Multilink procedure
- Charging information
- Direct call
- Hunt group
- On-line facility registration
- D-bit modification
- Local charging prevention
- Call redirection
- Network user identification
- Extended framer sequence numbering
- RPOA selection

Optional Call User Facilities

- Closed user group selection
- Bilateral closed user group selection
- Reverse charging
- RPOA selection
- Flow control parameter negotiation
- Fast select
- Throughput class negotiation
- Abbreviated address calling
- Charging information
- Transit delay selection and indication
- Call redirection notification
- Called line address modified notification
- Network user identification
- Closed user group with outgoing access selection

Table 2. X.25 packet network options.

Managing X.25 Packet Switched Networks

Vendor	Product
Amdahl Corp.	NMS Series 200: A combined T1, SNA, and X.25 management system based on a Sun workstation for monitoring, alarms, configuration, and billing. Supports Informix. Cost: \$25,000+
Amnet, Inc.	Nucleus 7000 Series: An X.25 monitor for Amnet networks. Full featured. Cost: \$40,000-45,000
AT&T	Accumaster/Integrator: A combined network management tool for all kinds of networks, including X.25. Based on AT&T 3B2. Full featured. Supports Informix. Cost: \$250,000
BBN Comm. Corp.	C/7 NOC; T700/CSM: A Digital Microvax-based, full-X.25 network management system. Supports Ingres. Cost: \$75,000+
Codex Corp.	9800: A multipurpose network management system for X.25 and other networks. Workstation based. Full featured. Cost: \$45,000+
CSC Infonet	NCC/PC: A PC-based X.25 management tool for users of Infonet.
Dynatech Comm., Inc.	Prism: A PC-based, full-featured X.25 network manager. Cost: \$10,000-20,000
Hughes Network Systems, Inc.	Graphics Network Operator Console: A PC-based network management tool for Hughes IPN packet switches.
IBM	NetView/PC: PC-based software that supports X.25 links in an SNA network. Requires VTAM. Cost: \$2,000-7,000
Infotron: Sys. Corp.	Integrated Network Manager: A combined T1-X.25 network management tool. Cost: \$35-82,000
Racal-Milgo	CMS Multi-link Four: A PC-based multipurpose X.25 network management system. Supports Oracle. Cost: \$3-10,000
Telenet Comm. Corp.	TP5000 NCC: A Prime-based, full-featured X.25 NCC for Telenet networks. Uses NMIS software. Supports Oracle.
Timeplex, Inc.	Time/View: A Sun workstation, combined T1, mux, and X.25 network control system. Supports Informix.
Tymnet	Supervisor: A full-featured, Sun workstation-based X.25 network manager. Supports Sybase.

Table 3. X.25 network management products.

proper network operations. The network control staff needs to understand both Lap B and Packet structures to troubleshoot the software intelligently.

The testing of X.25 software is similar to protocol testing of any major protocols: emulate the protocol (X.25), send out a message, and verify the correct response.

SECURITY AND CONFIGURATION MANAGEMENT

Users require authorization codes and security provisions, similar to those for VAN access. User IDs and passwords must also be assigned, just as in the VAN case. The network manager must select all the configuration parameters when the private network is set up. Fortunately, these are well defined and explained in the network setup instructions. Table 2 lists some of the features and options, called "facilities." Some facilities are mandatory, while others are truly optional. It is imperative for the network manager to consult with the PAD equipment vendor to determine which options are provided and to properly set these options for the user's particular network needs.

The network manager must choose which protocols to support, based on the host requirements in the corporate network. Today's packet switched networks support most of the common host access methods, which range from asynchronous to SDLC network protocols.

PERFORMANCE MANAGEMENT AND STAFFING REQUIREMENTS

The network manager must also monitor network performance, not only to detect circuit congestion, but also to ensure proper nodal operations. User statistics are collected at the nodes and returned to the NCC to aid in network support.

Network managers must also address organizational issues when a private packet network is deployed. Both hardware and software staff are required to properly support the network and its users. While a single engineering staff can support both packet switching and other architectures, it is wise to have a separate functional network control staff. The personnel needed to staff the X.25 network is quite similar to the sizing required for a conventional circuit switched network.

Managing X.25 Packet Switched Networks

Vendor	Product
Dynatech	Simon V: A full-featured test set with storage for protocol testing of X.25 lines as well as SNA and async circuits. Cost: \$15,000
Frederick Engineering	Feline-LT: A PC-based line monitor for X.25 and other protocols. Cost: \$2,000
Hewlett-Packard Co.	4954A: Protocol analyzer for X.21, X.25, X.75 and others.
Idacom Electr. Ltd.	IPT: A full-featured test set for ISDN testing as well as X.25 and X.75. Cost: \$25-35,000
Navtel, Inc.	DataTest Remote: Portable X.25 and SNA test set. Cost: \$4,000
Spectron	Datascope: A full-featured test set with storage for protocol testing of X.25 and other lines. Cost: \$10,000
Telebyte Technology	Netscope: Analyzer for X.25 level three only. Cost: \$1,800

Table 4. X.25 test equipment.

One difference, however, is that it is more critical for the staff to understand the X.25 software issues than software in a conventional network, because X.25 software is so complex and such an important part of the architecture.

COST CONTROL

Since the private network usually offsets the cost of making long-distance phone calls to reach the necessary hosts, a chargeback scheme, similar to that which might be provided for users of VANs, may be deployed. Such a scheme assists the network manager in controlling network usage; it makes users aware of the amount of time spent on the network.

Network managers can also employ usage monitoring to detect unnecessary use—this can help reduce the need to add nodes or upgrade line speeds and thereby achieve increased cost control.

Another way to minimize growth costs is to set up a hybrid network by linking a private network into a public VAN. In this way, low-volume users can access the hosts on the private network through a public VAN gateway.

X.25 NETWORK MANAGEMENT PRODUCTS

NETWORK MANAGEMENT SYSTEMS (NMSs)

Network management systems (NMSs) are particularly critical for private packet switched network installations. Each packet-switching equipment vendor offers such a system (see Table 3 for a list of vendors). Increasingly PC or workstation based, these systems are invaluable both for network monitoring and fault diagnosis. Each of the systems described in Table 3 works primarily with the individual vendor's packet-switching nodes, although many of the vendors are promising OSI compliance in the future, which should provide some degree of interoperability.

X.25 NETWORK MANAGEMENT TEST EQUIPMENT

There are many manufacturers of X.25 test equipment. There are a number of full-featured central-site testing units available, as well as remote testing products in small portable configurations. A list of some of the vendors and their products is provided in Table 4. □

Telecommunications Management Software

This report will help you to:

- Compare and contrast various approaches to telecommunications management.
 - Choose the approach that best fits your organizational needs.
 - Evaluate the reliability of potential vendors.
-
-

THE MICROCOMPUTER SOLUTION

Microcomputers may be used to collect, poll, and process call records and to provide reports on demand. Many other benefits of microcomputer-based telecommunications management deserve attention (Table 1).

Inexpensive floppy disks can be used to store call records conveniently for later reprocessing, historical inquiry, and traffic analysis. Floppy disks are much simpler to use than the magnetic tapes commonly used to store data for mainframe processing. Menu-driven microcomputer products provide the user with choices that quickly retrieve report and perform database maintenance, allowing users with little or no experience to be instantly productive. Some microcomputer software programs include features that allow users to customize telecommunications reports.

The same hardware can be used for call accounting and for other telecommunications management jobs, such as database and telephone directory maintenance, polling, inventory control, service order tracking, traffic analysis, and network optimization. In fact, sophisticated windowing programs and other background-mode techniques now allow the microcomputer to record call data while the operator performs other tasks. For example, while calls are being

recorded, the operator can use another window to access an electronic telephone directory, check the status of a service order, or run a spreadsheet on the telecommunications budget. Such multitasking capabilities provide telecommunications managers with additional justification for purchasing standalone gear and to preempt management concerns about dedicating a microcomputer to call-record collection. Finally, hardware support for microcomputers is widely available and, for the most part, reliable.

But the apparent convenience and user-friendliness of microcomputer-based telecommunications management often obscure its inherent limitations. Despite the impressive collection and storage capacities of some microcomputers—as much as one million call records a month—what really counts is the number of call records that can be processed at one time and how long it takes to format and run the final reports. A microcomputer that can collect up to 500,000 call records a month from 2,200 extensions, for example, may produce only summary reports, or retrieve and print call data one station at a time. This may not be adequate for running detailed and consolidated usage reports on a timely basis. Depending on the number and complexity of reports required, and the rating scheme used to cost each call, processing that many call records can consume as much as 120 worker-hours per month, requiring the supervision of a full-time operator.

Many vendors unintentionally confuse prospective customers when they use the term “processing.” To equate it with “call collection” alone is misleading.

This Datapro report is based on “Surveying the Software that Manages Telecommunications,” by Nathan J. Muller, From *Data Communications* magazine, March 1987, pp. 201-202, 205-206, 209, and 211. © 1987 by McGraw-Hill, Inc. Used by permission.

Telecommunications Management Software

ADVANTAGES	DISADVANTAGES
MULTIFUNCTIONAL: COLLECTION POLLING REPORT PROCESSING DATABASE MAINTENANCE	LIMITED MEMORY AND PROCESSING POWER
MULTITASKING CAPABILITY	
REPORTS ON DEMAND	
INEXPENSIVE	LIMITED CALL-RATING
CONVENIENT USER FRIENDLY	REQUIRES USER TO BE RESPONSIBLE FOR TURNKEY OPERATION LIMITED REPORTING FLEXIBILITY
HARDWARE SUPPORT READILY AVAILABLE; ECONOMICAL FOR SINGLE-SITE TELECOMMUNICATIONS NETWORKS WITH LESS THAN 500 EXTENSIONS	

Table 1. Summary of microcomputer processing.

Processing actually refers to the systematic execution of operations on a call record to arrive at a desired result—that is, calls priced approximately and arranged in a meaningful report format. Call collection, then, is only the first step in a much more complicated process that leads to report generation. When vendors say that their microcomputer-based machines can “process” or “handle” 50,000 calls an hour, what they are really trying to say is that when their microcomputers function as call collectors, they can collect and store 50,000 raw call records.

Generally, microcomputers cannot rate calls according to appropriate tariffs for intra- and inter-state or intra- and interLATA (Local Access and Transport Areas) calls. The same holds for calls to Canada, Mexico, Hawaii, Alaska, and international (011) locations. With their limited memory, storage, and computing capabilities, microcomputers must approximate rather than accurately rate calls using V&H (vertical and horizontal) tables. These tables provide, among other things, exchange coordinates that are used for calculating the distance of calls. But this method is only 90 percent to 95 percent as accurate as rating calls with actual tariffs, which may be close enough even for high-volume users.

Using the V&H costing method, Xtend Communications Corp. in New York, for example, can collect and

price up to 30,000 calls per hour in real time with a microcomputer; it can have some reports ready in only 30 seconds. (This kind of performance is more the exception than the rule.)

Microcomputers virtually require that users take full responsibility for quality control. It is often a full-time job to monitor equipment and keep it running smoothly. Although microcomputer software costs much less than that for mainframes, the limitations of most microcomputers should be carefully weighed against the promise of short-term savings.

If user call-costing and -reporting needs are relatively simple, the microcomputer processing option may be the best choice. This alternative also gives the added benefits of convenience and multiple functions. As needs change, and insight and expertise in telecommunications management grow, software modules and memory capacity can be added at only an incremental cost. The microprocessing option is also an effective way for telecommunications managers to end their dependence on the MIS/DP (management information systems/data processing) shop and retain total control of their corporate mission. From top management's perspective, the microprocessing option should stop the finger-pointing between MIS/DP and telecommunications.

Alternatively, users can migrate to minicomputer or mainframe products, perhaps with the same vendor. When the time comes to make that decision, the user will have built up the knowledge base from which to make more informed choices.

MINICOMPUTER PROCESSING

Minicomputers constitute another on-site processing option. For many years, this method was used almost exclusively by the hotel and motel industry for on-demand billing. Vendors developed specialized hardware and software packages for this lucrative market niche. But as the market became saturated, many of these vendors enhanced their offerings for broader use. Today, approximately 150 vendors compete with about 500 software offerings in a market that will approach \$225 million in revenue by the end of 1988.

Because today's microcomputers can provide nearly the same storage capacity and processing power as medium-range minicomputers, the market for minicomputer-based telecommunications management products is not expanding. Advances in microcomputers have come faster than advances in minicomputers and mainframes. However, many mainframe software vendors also offer their products for use on minicomputers.

Telecommunications Management Software

For the most part, the following discussion of mainframe processing options also applies to minicomputer telecommunications management packages.

MAINFRAME SOLUTIONS

Mainframe processing comes in two versions: service bureau and license agreement.

With a service bureau, the user provides the vendor with call detail records through magnetic tape or floppy disk, or through electronic transmission over telephone lines under a polling arrangement. The vendor uses its mainframe computer to process the data and to generate the reports requested by the customer. A reasonable turnaround time for this kind of report processing is one week.

Under a license arrangement, the user typically signs an agreement with the vendor. The document authorizes use of the software at one or more sites where the user's hardware is located. Such a license agreement is common practice among vendors because it allows them to exercise more control over the proprietary nature of their software than copyright laws usually afford with an outright purchase. Under a license agreement, the user is obligated to maintain confidentiality and adhere to provisions in the agreement governing the use and disclosure of the software product.

The licensed software package usually includes tariff tables required for accurate call rating and a database of V&H coordinates for assigning city and state locations to long-distance calls made over bulk-rated facilities. It also includes a database defining the characteristics of the user's telecommunications setup that will use this data to generate the desired reports.

SITE LICENSE VERSUS SERVICE BUREAU

Choosing between the license or service bureau may hinge on internal capabilities, specifically the amount of time and effort the MIS/DP shop wishes to expend in becoming experts at telecommunications management (Tables 2 and 3). Installing new tariff tables, updating the database to reflect changes in the network, and verifying input data and report runs are all time-consuming chores that require expertise and staff continuity. Ultimately, it may be easier and more economical to use a reliable service bureau for such tasks. A good service bureau also provides assistance in interpreting the reports.

Business with a large number of extensions or account codes, coupled with constant equipment moves

and changes, should ask if the vendor offers a front-end capability for entering and uploading this data to the mainframe. Menu-driven microcomputer data entry permits users to easily record, review, and edit the database—a more convenient and accurate process than filling out forms for keying by the vendor's data-entry staff.

Remember that with multinode networks that use private branch exchanges from different manufacturers, each PBX type formats call records differently. Trying to handle a variety of formats on an in-house mainframe may take up valuable processing time and cause program maintenance nightmares for the MIS/DP staff.

Distinct disadvantages arise when a telecommunications management application uses a data processing center that relies heavily on magnetic tape. Since magnetic tape is not pollable and since different PBX types use proprietary call record formats, the data stored on each magnetic tape must be preprocessed into a common format before mainframe report processing can take place. Depending on the number of PBX types the network has, this could turn into a cumbersome procedure that wastes mainframe resources and increases the chance for human error.

Magnetic tape drives are mechanical devices that require scheduled preventive maintenance. Calls are not recorded during maintenance downtime unless redundant magnetic tape hardware is operating—a very expensive proposition.

ADVANTAGES	DISADVANTAGES
VIRTUALLY UNLIMITED MEMORY AND PROCESS-POWER	MAGNETIC TAPE MAY PROVE CUMBERSOME AND TAX MAINFRAME RESOURCES
PRECISE CALL RATING	HIGHER UP-FRONT COSTS FOR SOFTWARE
INCREASED REPORTING FLEXIBILITY	REQUIRES KNOWLEDGEABLE DATA PROCESSING PERSONNEL; MINIMUM TURNOVER
TIMELY REPORT PROCESSING	REQUIRES USER TO TAKE CHARGE OF TURNKEY OPERATION
ECONOMICAL FOR MULTI-NODE USERS WITH MORE THAN 500 EXTENSIONS	

Table 2. Summary of minicomputer/mainframe processing (license arrangement).

Telecommunications Management Software

ADVANTAGES	DISADVANTAGES
VIRTUALLY UNLIMITED MEMORY AND PROCESSING POWER SUPPLIED BY VENDOR	DELIVERING DATA THROUGH MAGNETIC TAPE OR FLOPPY DISK IS RISKY
PRECISE CALL RATING	NOT ECONOMICAL FOR RAPIDLY GROWING COMPANIES WITH CONSTANT EQUIPMENT MOVES AND CHANGES
NO CAPITAL INVESTMENT; ONLY MONTHLY PAYMENTS FOR THE SERVICE	
INCREASED REPORTING FLEXIBILITY	LONG-TERM VIABILITY OF SERVICE HINGES ON CONTINUED SUCCESS OF VENDOR
REASONABLE TURNAROUND TIME FOR REPORTS	
ECONOMICAL FOR TELECOMMUNICATIONS NETWORKS WITH UP TO 1,000 EXTENSIONS AND LIMITED DP RESOURCES	
CUSTOMER SERVICE AVAILABLE ON CONTINUING BASIS	
QUALITY CONTROL IS ASSURED	

Table 3. Summary of minicomputer/mainframe processing (service bureau).

Furthermore, it is not always easy to spot a malfunctioning magnetic tape drive; a misaligned head, for example, can cause call-recording problems. Many times such a malfunction is not discovered until a technician stumbles on it during a preventive maintenance visit—or when the MIS/DP manager notifies the telecommunications administrator that a blank tape was sent to the data processing center. Even more basic, it is easy for inexperienced operators to mistakenly install a magnetic tape on the drive unit. This can result in lost call records, wasted computer time, and delayed report processing. Moreover, magnetic tape reels can be lost in transit, damaged from mishandling, or destroyed in an accident.

Magnetic tapes require careful administration and security for keeping track of spares, backups, tapes in transit, tapes awaiting processing, and blanks ready to be used. Not only does this burden the user with unnecessary overhead costs, but one foul-up can throw a tightly scheduled data processing operation

into chaos. The alternative—submitting finished reports late—may result in stale information that is of limited use to the telecommunications manager.

Pollable solid-state collection devices, black boxes that hang off PBXs and are responsible only for collecting call records, eliminate these problems. They are relatively inexpensive: from less than \$3,000 for a store-and-forward device like Account-a-Call's Tadpoll to \$6,000 for Sarasota, FL-based ComDev's intelligent STU-3B, which is available only through distributors. These devices install easily and quickly without PBX modifications or service interruptions. Because they can be polled, they also provide substantial cost savings in labor, magnetic tape shipping, and mainframe processing.

If reports are processed on an in-house mainframe and the network involved is composed of different PBX types, these store-and-forward devices do nothing to simplify the task of format translation. Although they operate unattended and store as many as 60,000 call records, they do not have the intelligence necessary to convert this data into a common call record format, a chore that will ultimately tax mainframe resources. However, for recording needs on a homogeneous network, these low-cost call data collectors perform this simple task reliably.

Alternatively, look for an intelligent Station Message Detail Recorder (SMDR) that will process call records from virtually any PBX into a common format. Although they cost more, these smart devices make data collection and report processing much easier, and they provide flexibility in choosing additional PBXs for network expansion. They can be purchased without regard for PBX type (see Figure 1). An additional benefit is that a single device can accept data from multiple PBXs simultaneously. For example, ComDev's STU-3B collects data from up to six PBXs simultaneously, while Telco Research's TRU Recorder records data from up to 15 PBXs simultaneously. If saving money on gear is an objective, look for a service bureau that will poll inexpensive store-and-forward data collectors and assemble the diverse call record data in a common format for call costing and report processing.

In addition to eliminating the risk, labor-intensiveness, and cost of magnetic tape delivery, the polling process can enhance data integrity by auditing the data from the point or origin to the point of delivery. Blocks of data are sequentially numbered at the point of origin and are "checkpointed" as they are teleprocessed to the destination polling equipment. In this cyclic redundancy check, the number of data blocks is verified by both the sending and receiving units. Sometimes the ACK/NAK (acknowledgment/

Telecommunications Management Software

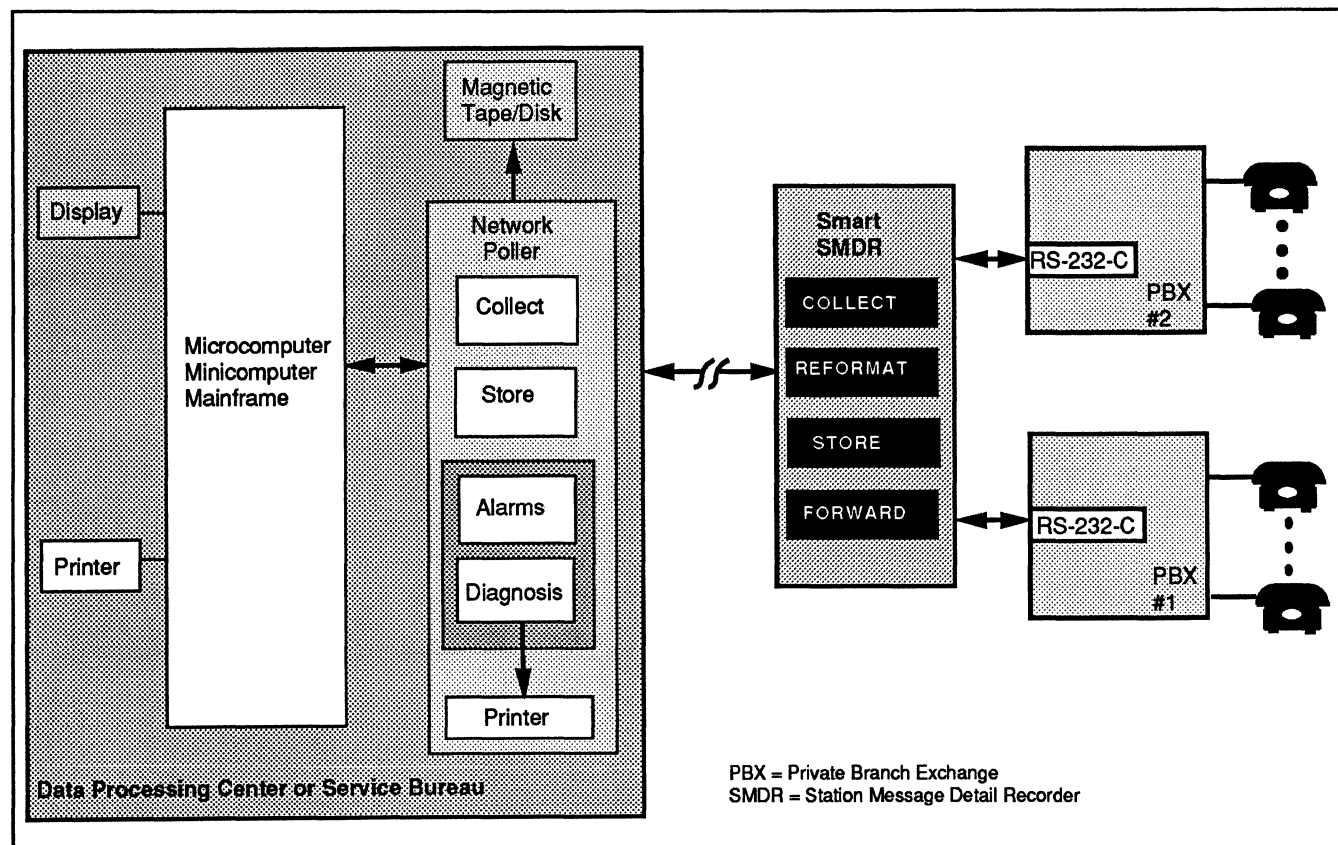


Figure 1. Multinode teleprocessing. SMDRs can collect data from diverse PBXs, formatting it for use by a service bureau or on-site data processing center. SMDR buyers need not consider the type of PBX sending data. Although up-front costs are greater than with store-and-forward devices, overall network management costs may be less.

negative acknowledgment) process is used to confirm the receipt of data. In the event that incomplete or damaged data arrives at the network poller, repolls are initiated automatically until all data blocks are accounted for and arrive at the destination error-free.

Polling provides a degree of flexibility to data collection in that the process may be initiated as demand warrants or according to a routine late-night schedule that takes advantages of lower transmission costs. And most pollers are designed for unattended operation. If users have high call volume and plan to use their microcomputers in a multitasking mode, polling must be done more frequently because there is less memory available for storage.

But even the selection of polling equipment requires careful evaluation:

- In the event of a power outage, will the poller turn off and stay off, or will it restart automatically and initiate a repoll when the power resumes?

- If there are different SMDR types on a multinode network, will the poller adjust automatically to the error-checking protocol of the SMDR device being polled?
- Does the poller have enough intelligence to switch lines automatically or to hang up and dial again when it encounters poor line quality?
- Does the poller allow remote activation for polling on demand?
- Does the poller furnish a complete report of its activities?
- Does it have the capability to perform remote diagnostics on PBXs as well as on the SMDR devices?

Answers to these questions are critical to the long-term reliability of data collection and, ultimately, to the viability of in-house report processing on multi-node networks. When these questions are applied to most microcomputer-based pollers, there may be trade-offs between reliability and convenience and low cost. Unlike their minicomputer or mainframe

Telecommunications Management Software

equivalents, microcomputers may not have the intelligence necessary to automatically reinstate polling after a power outage.

The comprehensiveness of tariff used by the vendor ultimately has an impact on the accuracy of cost allocation. With the implementation of LATAs in 1984, there are now many different rate structures for inter- and intrastate calls, and others for inter- or intra-LATA calls. With 164 LATAs defining the service areas of the major telephone companies and another 25 for independent telephone companies, the task of rating calls is complex. Thus, find out how the vendor developed and implemented its new rating package when the LATA schema went into effect after the divestiture of AT&T. Also, find out if its customers experienced any problems related to the change-over.

Answers to these questions may determine the reliability of vendor service and indicate future responses to changes in the operating environment.

Other points to keep in mind include:

- The advantages of being able to rate calls differently from department to department.
- In addition to costing calls by tariff, costing may be user defined as a flat cost-per-call, cost-per-minute, or cost-per-minute with differential rates for the first and each additional minute.
- Calls placed over bulk facilities may be costed at the equivalent DDD (direct distance dialing) rate or calculated as a percentage of the total corporate telecommunications cost.
- Time-of-day costing allows the application of appropriate evening, weekend, and holiday discounts to DDD calls and calls made over alternative interchange facilities.

When used in conjunction with account codes, any of these costing schemes can be used for client billing. Some software vendors are adept at providing all of these rating options and will even work closely with users to determine their exact requirements.

A thorough assessment of organizational needs will also determine the importance of having a vendor that can accurately rate calls originating and terminating in the same state. Calls originated and terminated in California are rated differently from calls originated and terminated in New York. Consequently, if a company has multiple locations, extra tariff files have to be maintained to rate intrastate calls. Although this information is readily available

from such reputable organizations as CCMI/McGraw-Hill, some vendors are not properly equipped for such an undertaking and will try to make a case that other, less accurate call rating schemes, such as flat costs per minute, may also work.

NETWORK OPTIMIZATION

Post-divestiture competition has resulted in more frequent and more complex tariff changes, making it almost impossible for telecommunications managers to evaluate everything available to ensure an optimally configured network. If a voice network has 500 or more extensions and relies extensively on alternative carriers and bulk-billed services, a network optimization study should be performed at least annually, or more frequently as changes in available services, equipment, and network configuration dictate.

Network optimization studies yield valuable information that can be used for the following purposes:

- Ensure maximum savings with alternative carriers by finding out how they serve specific calling patterns.
- Justify FX, Tie Line, or WATS services with accurate and comprehensive information useful in making decisions.
- Determine through "what-if" analysis the potential impact of proposed tariff changes on the long-distance bills.
- Validate network reconfiguration proposals for top management review and approval.
- Set the pace for corporate fiscal responsibility.
- Sharpen the company's competitive edge with an efficient telecommunications network.

Paradoxically, vendors and consulting firms specializing in network optimization exclusively are severely handicapped compared with those that offer call accounting on a license and service bureau basis. While on first inspection the network optimization companies appear to be less biased, the second group of vendors have the advantage of already maintaining tariffs and huge databases, endowed with the computing power necessary to provide comprehensive call accounting services. Generally, they are going to be more reliable and thorough when it comes to providing network optimization services.

Telecommunications Management Software

In addition, call accounting firms can accommodate the data from a number of different PABX types. Generally, the turnaround time for a network optimization study will be shorter than with vendors that lack these capabilities because call-accounting firms already know the correct call record format used on each PBX and can read the data without encountering translation problems.

Even if network optimization is only a remote consideration, users may want to choose a call-accounting vendor that has a reliable track record with such services. The charges for network optimization can be as much as 20 percent less because database setup has already been done to provide call accounting. A reasonable turnaround time for a service bureau network optimization study is four to six weeks, depending on the number of network nodes plus the nature and scope of transmission facilities to be considered.

But caution should be taken with network optimization vendors who promise a certain percentage of savings or base their fees on a percentage of savings. Usually, such vendors will not allow much input about the type of long-distance and bulk-billed services to be included in the study. These vendors want to make such decisions themselves because they can then deliver the greatest amount of savings and thus higher fees for themselves.

Also, be on guard with a vendor who happens to be a subsidiary of a telephone company. Impartiality may take a backseat to achieving larger corporate objectives.

If total control over the network optimization process is desired, and the extra time and effort to become proficient in network planning are available, there are microcomputer-based traffic analysis and network optimization packages available. Such packages include tariff and availability information for all major long-distance carriers. Some packages include a database of services available on an exchange-by-exchange basis, especially useful for comparing calls to cities that may not be on different discount carriers, and so will cost more than on-network calls. Also included in these packages is a program for analyzing real-time calling patterns, which can be used to select the most appropriate long-distance services.

Although such packages are effective for optimizing a single location, performing optimization studies on each node of a large nationwide network can be cumbersome and time consuming, requiring dedicated staff. In this case, the mainframe processing capabilities of a service bureau might be more efficient.

VENDOR RELIABILITY

Be wary about using the service bureaus of PBX manufacturers or equipment vendors. Being hardware-oriented, their call-accounting capabilities are usually limited to their own brand of PBX, which could limit equipment choices when it comes time to expand the network. When software problems arise, hardware vendors can be very slow with problem resolution.

Find out if telecommunications management is a major or minor part of the vendor's business. If it is only a sideline, the pressing needs of call-accounting customers may not receive top priority treatment.

Before committing to a vendor:

- Insist on reviewing product technical documentation for scope and clarity—and to verify the claims of salespeople. If the in-house MS/DP shop has problems attracting and keeping qualified staff, the quality of technical documentation may prove to be critically important if the choice is made to license mainframe software.
- Consider the ability of the vendor to customize software to meet unique organizational requirements.
- When calling references, be sure to ask about the timeliness with which customization was completed, the cooperativeness of the vendor in ironing out bugs in the program, and whether the final product matched the buyer's expectations.
- If a reference who no longer uses the vendor's products appears, find out why the company changed vendors.
- If a vendor appears to be giving away too much to get business, be suspicious. Getting telecommunications management at a bargain price will do no good if the vendor goes out of business and leaves no ongoing support. The call-accounting industry is extremely volatile. Some of the best-known companies are only marginally profitable; others have gone under and have resurfaced as new entities with scaled-down product offerings and severely limited levels of customer service.
- Don't accept standard contracts. All provisions of the contract are negotiable, including price and terms. As in any other business transaction, the seller sets the price and the buyer sets the terms. Just because a contract is typeset doesn't mean users can't spell out the term in as much detail as they deem necessary.

Telecommunications Management Software

- If custom software is part of the purchase, throw in a penalty clause for late delivery or a "weasel" clause that delivers a full refund if the product doesn't perform as promised after a reasonable prove-in period.
- To guard against the loss of technical support, make sure that the purchase includes a program source code in the event that the vendor goes belly-up or discontinues the product. The source code should be deliverable automatically from an escrow account or from a third party specializing in such services. Be advised that this provision in the contract requires the assistance of an attorney who is experienced in matters of software protection—it's too easy to overturn these provisions in court. Also, make certain that whenever the product is updated, the source code in escrow is also updated.

If vendors balk at any of these notions, hand them a list of their competitors who are eager to get the business.

LOOK BEFORE YOU LEAP

Whether a microcomputer, minicomputer, or mainframe call-accounting package, it carries a commitment to a long-term relationship and a dependency on the vendor for tariff updates, software enhancements, customer service, and so on. Vendors must make the case that they will be around to provide support for many years to come. It is the user's responsibility to exercise due diligence with respect to vendor selection by obtaining customer references, as well as credit and financial information. An added measure of caution is called for when dealing with privately held firms. In this industry it is these companies that have been most prone to mismanagement—falling behind in product development and having serious problems with cash flow.

If the decision is made to rely on a consultant for advice, show that individual no mercy; be alert to potential under-the-table relationships with vendors. Some consultants are paid retainer fees by vendors, given as an incentive to bring in new business. Insist on a thorough justification for all decisions. □

The Telemanagement Symphony

This report will help you to:

- Evaluate the various types of telecommunications system management software available.
 - Plan for the implementation of a telemanagement system.
 - Anticipate future trends in telemanagement.
-
-

Symphony conductors and telecommunications managers have more in common than they realize.

Just as the conductor uses a baton to give the musicians instructions and “manage” the concert, so do organizations use telemanagement software to control their networks and manage their telecommunications environment.

And just as the conductor must coordinate the harmonious interaction of such diverse types of instruments as woodwinds, strings and percussion, so too must a telecommunications manager direct the different factions within the company to work together to choose the telemanagement system best suited to the organization’s needs.

While the initial choice of operating environment may seem to be driven by corporate data processing standards or cost investment differences, users should assess the political climate between the telecommunications and DP/MIS groups before reaching a decision. Often, corporate politics intervene in a telemanagement system’s success.

This Datapro report is based on “The Telemanagement Symphony,” by Lillian Goleniewski and Andrea Wells, The Lido Organization, Inc., from *Network World*, pp. 32-35, 42, February 1989. © 1989, NW Publishing, Inc. Reprinted by permission.

POLITICAL CONCERNS

After defining the basic system capacity requirements necessary to support their telemanagement applications, users should ask themselves the following questions:

- Is MIS merged with telecommunications, and do the groups work well together? Many a mainframe system falls short of the mark because DP doesn’t understand telecommunications and the important role telemanagement systems play.
- Does the telecommunications department want autonomous control of the telemanagement functions? If so, a departmental system, such as a dedicated minicomputer, may offer the most benefits. However, with autonomy comes more responsibility for the system—a system that will become increasingly crucial as its benefits become more visible to executive management.
- Will the use of a telemanagement system increase the political clout of the group implementing it? The clever use of information systems will enhance the bottom line of any organization. This opportunity to command boardroom visibility may influence the user’s choice of hardware operating environments.
- Who will authorize the purchase of a telemanagement system, and how critical is price to this

The Telemanagement Symphony

choice? How will politics and required purchasing approval affect the time involved in approving and acquiring a product?

Keeping in mind the four political questions just posed, we can look at today's telemanagement alternatives: the mainframe, minicomputer, microcomputer and service bureau offerings.

Of course, the most critical aspect of system selection is how the applications provided fit with internal information needs. Beyond system capability, users must always remember to assess the more subtle political concerns of system control, internal cooperation, cost sensitivity, and managerial clout.

Each of these choices has pros and cons, and many of the traditional distinctions between operating environments and the types of telemanagement functionality available on each are disappearing. Certain user-driven needs, such as on-line inquiry, multitasking, and historical reporting, are becoming more and more common in all types of software packages. Some of the most successful installations have combined operating systems to maximize telemanagement benefits.

MAINFRAME-BASED SOFTWARE

The use of a mainframe computer may be dictated by virtue of its existence in the company, or it might be the favored environment for complex systems that require sophisticated MIS support. Mainframes can handle the batch processing of high-volume call-accounting systems while providing fully integrated telemanagement applications with sophisticated functions for large organizations and communications networks. (See Table 1.)

If centralized management and control is required, mainframes can be used as the focal data base point, working in a distributed DP mode to provide individual locations with some level of telemanagement functionality through the use of microcomputers. The remote locations can then interact with the central data base to maintain a corporate network profile.

Mainframe telemanagement systems have many advantages. Not only do they have more processing power and higher speeds, they also have virtually no capacity limitations. They allow for multiple users and can run a variety of sophisticated applications. Generally, MIS departments are familiar with mainframe computing and will have little or no trouble supporting it—meaning less support responsibility for the telecommunications staff.

Mainframe telemanagement systems have drawbacks, however. The telecommunications department has less direct control over the system, and consequent job scheduling, and will have to interact with MIS more. Although MIS is familiar with the mainframe environment, it has less operational expertise when dealing with telemanagement system support and maintenance. Users will have to put up with slower responses to change requests. They will also find higher costs associated with mainframe usage for the initial software purchase, and, if the organization practices CPU chargeback, they will have higher internal costs.

Top industry providers of mainframe telemanagement systems include: Telco Research Corp. of Nashville; Stonehouse and Co. of Dallas; Communications Design Corp. of Stamford, CT; and Cincinnati Bell Information Systems' Communications Management Systems of McLean, VA.

MINI-BASED SOFTWARE

While the minicomputer telemanagement software option has commonly been used as a departmental solution that provides applications as sophisticated as the mainframe, the traditional distinctions between minicomputers and supermicrocomputers (32-bit processors with windowing capabilities and high-resolution graphics, such as Sun Microsystems, Inc. workstations) are becoming more and more hazy. For organizations with systems requiring multiple users, multiple applications, and high-volume capability, as well as those requiring multitasking and system partitioning, the minicomputer presents a viable alternative.

Using minicomputers for telemanagement allows users to have departmental autonomy, as well as immediate control over report and request turnaround times. The minicomputer operating environment can also accommodate multitasking and multiuser capabilities, and it has lower processing costs than the mainframe. Recent cost reductions in the minicomputer market have also enhanced the practicality of the minicomputer telemanagement solution. Both Hewlett-Packard Co. and AT&T offer minicomputers priced at less than \$15,000.

The disadvantages of minicomputer software include possible limits on system capacity, greater capital investment than a microcomputer solution, and a requirement for more telecommunications expertise and training in administering an autonomous system as opposed to an MIS system.

The Telemangement Symphony

	Mainframe	Minicomputer	Microcomputer	Service bureau
Advantages	<ul style="list-style-type: none"> • Maximum capacity limits • Multiple users • Processing power and speed • Less telecommunications support responsibility • System partitioning • Experienced mainframe support and system backup 	<ul style="list-style-type: none"> • High capacity levels • Multiple users • Dedicated departmental processor • System partitioning 	<ul style="list-style-type: none"> • Less capital expense and maintenance expense than mainframe or minicomputer • Dedicated departmental processor • Full system autonomy • User-friendly • Most common DOS interfaces 	<ul style="list-style-type: none"> • Unlimited capacity • High level of telemangement system expertise • Off-load telemangement system responsibility • No capital expense
Disadvantages	<ul style="list-style-type: none"> • Less direct system control • MIS less responsive to change requests • Less MIS familiarity with telemangement • Higher purchase and support costs • Higher operational expertise necessary 	<ul style="list-style-type: none"> • Requires more operational expertise from telecommunications staff • Limits on system capacity • Minicomputer market challenged by new microcomputer capabilities • Higher capital expense than microcomputer or service bureau 	<ul style="list-style-type: none"> • Number of users limited by system configuration • Limits on system capacity • Slower processing and printing • Most critical backup procedures necessary 	<ul style="list-style-type: none"> • Report lag times • User rarely has on-line capability • Applications generally limited • Data security hampered
Political Profile	<ul style="list-style-type: none"> • Mainframe telemangement system requires a strong level of interdepartmental cooperation, more technical expertise and financial investment 	<ul style="list-style-type: none"> • Departmental system control still requires sophisticated computer and software support • Cost varies with both hardware and software 	<ul style="list-style-type: none"> • Direct telecommunications control over telemangement system at the microcomputer level. Telecommunications will have full system responsibility. • Most responsive to price-sensitive market 	<ul style="list-style-type: none"> • Users have no control over system administration and are totally dependent on vendor for service. • Recurring costs

Source: The Lido Organization., Inc., Mill Valley, CA.

Table 1. Telemangement software options.

The price range for minicomputer software packages is very wide. They can be found at prices comparable to both high-end microcomputers and low-end to mid-range mainframe packages.

The major vendors of minicomputer-based software are: Communications Group, Inc. of King of Prussia, PA; Comsoft Management Systems, Inc. of Houston; NEC Information Systems, Inc. of Boxborough, MA; Telco Research; and TelWatch, Inc. of Boulder, CO.

MICRO-BASED SOFTWARE

The microcomputer market for telemangement software is growing quickly. The introduction of the more powerful Intel Corp. 80386 processing chips and local networks, as well as the promise of mul-

tiuser capabilities through IBM's OS/2, will remove many of the traditional limitations of microcomputer processing.

Today, call collection buffers and background/foreground operations make multitasking possible. Multiple users can access microcomputer telemangement systems today by operating in a local network environment, using multiuser personal computers and by operating in a UNIX environment. The predominant telemangement application served by traditional personal computer systems, such as the IBM Personal Computer XT and AT, remains single-function call accounting.

But new and more powerful microcomputers and increased user demand for complete, integrated telemangement have focused much attention on the development of expanded functionality for micro-

The Telemangement Symphony

computer-based systems. The workstation is evolving into the dominant user interface to telemangement: for example, AT&T's NetPartner product for Centrex, Software-Defined Network, and Integrated Services Digital Network management uses Sun workstations. This evolution has created a new generation of telecommunications and net management based on distributed DP approaches.

Microcomputer-based telemangement systems require less capital investment than those based on minicomputers or mainframes, and they have lower processing costs.

Users also have the advantages of more user-friendly interfaces and more direct control over report and request turnaround times. Another important advantage of a microcomputer-based system is that the users have total autonomy and direct control over the system.

However, there are several disadvantages to using microcomputer-based telemangement. System capacity and processing power are limited. With traditional personal computers, call-record costing and report-printing processes can be slow, depending on volume. Also, if not operating in a local network or UNIX environment, only one user can access the system at a time. Total telecommunications departmental responsibility means an increased need for internal computing knowledge and training.

Again, it is important to notice the extent of improvements possible as microcomputer capabilities expand. Quite a few systems now available provide real-time call costing, allowing faster call processing when ad hoc reporting is necessary. The increasing availability of on-line inquiry eliminates the need to generate full reports. Multiuser systems can be configured around a local network using UNIX and, in the future, OS/2.

The top vendors of integrated software applications designed to run on a microcomputer-based system are Communications Group; Comsoft Management Systems; Softcom, Inc. of New York; Telecommunications System Management of Harvester, MO; Tel-Watch; and Xtend Communications of New York.

New facilities management products that are being integrated with existing call-accounting software for microcomputers are announced almost monthly. New vendors in this market include Summa Four, Inc. of Manchester, NH, and Xiox Corp. of Burlingame, CA.

SERVICE BUREAUS

While the overall trend in telemangement is toward customer-owned software, the service bureau marketplace can still provide a viable solution. For either internal or resale billing, a service bureau can handle the responsibilities of managing systems. Also, off-loading management tasks to a service bureau will obviously eliminate the high capital expenses associated with an on-site telemangement system. Service bureau providers have also begun to sell on-site facilities management software as well as to provide off-site call-accounting services.

The advantages associated with service bureau applications include less internal responsibility for start-up and system cut-overs, and greater vendor familiarity and expertise with telemangement functions. Cost savings can also be realized by off-loading the need for internal staff time, training, and capital expense.

The disadvantages associated with the service bureau approaches center around the inability to directly access and manipulate the system applications. This can include slow turnaround times and reporting, lack of report feature customization often provided with custom report generators, threatened data security, and higher processing costs than with internal utilization.

As the end-user demands for telemangement system control and functionality have increased, service bureaus have looked less and less attractive. However, service bureau providers can be expected to expand the range of their services.

Better-known service bureau providers include Account-A-Call Corp. of Burbank, CA; Aud-Cyn Associates, Inc. of Parsippany, NJ; Communication Sciences, Inc. of Edison, NJ; Communications Group; Communications Management Systems, Inc. of McLean, VA; and Compco, Inc. of Brentwood, TN.

When beginning the search for a vendor of telemangement software, remember the nature of this particular industry within telecommunications. With the rapid growth of software packages on the market today, certain offerings will obviously be unable to compete for any extended length of time.

Many new vendors offer excellent telemangement products, but when actually choosing a provider, users should try to ascertain how stable the vendor will be a few years from now. Distributors of telemangement software may not be fully committed to supporting the products they are selling. Users must assess alternative support avenues for the future.

The Telemanagement Symphony

ANCILLARY HARDWARE

Besides choosing operating hardware, users must assess if and where they need peripherals hardware, as well as whether they want the peripheral hardware attached directly or remotely to their telemanagement systems.

Peripheral hardware captures and records data from a switch or switches, and then stores that data temporarily until it is polled by the system. Both call detail records and station message detail recording (SMDR) data and alarms can be captured and fed into the telemanagement system. SMDR data can be translated and then fed to the call-accounting and traffic modules for processing.

Alarms collected from a switch can be monitored at the site of collection and fed into a problem-tracking module to assist in pinpointing failures—failures that may direct attention to a deteriorating situation. This allows users to rectify a situation before it becomes a full-blown problem.

Many call-accounting products require a buffer box, which allows the user to perform other tasks with the telemanagement system and still ensure that data is collected from the switch at all times.

Alarm-monitoring devices for telemanagement systems are still quite new. Several firms provide switch-specific alarm monitoring and a basic translation of those alarms with integration into the problem-management module. Future applications of alarm monitoring could include some level of expert system that might receive the alarm, search its records for similar problems, and then actually open a trouble ticket with an advised resolution already provided.

Peripheral hardware is manufactured by several firms, including Com Dev, Inc. of Sarasota, FL; TSB International, Inc. of Rexdale, Ontario; and Western Telematic, Inc. of Santa Ana, CA. Many telemanagement system providers will buy products from OEMs, put their own names on them, and link them into their system.

When placed at a remote site to be polled by the telemanagement system or when used as a direct buffering device, the equipment should include several important features. LED capacity thresholds allow users to know how full the device's memory is. Error checking of data, battery backup in case of failures, format compatibility and translation, and internal modems also offer advantages. Each device must be defined for its specific use, either as a simple buffer, a

remote collection device, or a remote device that may include some reporting or printing capability at the remote location.

THE RIGHT OPERATING SYSTEM

Users must also choose the proper operating system for their needs. All computer systems use either one particular operating system or accommodate a variety. The final system configuration will indicate vendor options. But at the planning stage, users may still choose an operating system that both enhances their particular needs and points to preferred vendor selections.

Basically, the operating system is the supplied software or firmware that makes the hardware function and makes the computer's power available to the user. Essentially working as a resource manager, the operating system defines user interfaces, allocates shared resources among users, facilitates I/O, and permits users to recover from errors. This type of resource manager will then interface with system users, application programmers, software programs, and hardware.

Some common operating systems are: MS-DOS; PC-DOS for IBM Personal Computers and compatibles; UNIX for AT&T and HP computers, and NEC Corp.'s Micro XL; MVS and VM for IBM mainframes; VMS for the Digital Equipment Corp. VAX; and ITOS for NEC computers.

TELEMANAGEMENT TRENDS

Today's trends in technology—fourth-generation languages, portability, relational data bases, distributed data processing, local networks, and artificial intelligence—can be overwhelming. But they also offer the telemanagement user new functionality and control. The newest features of telemanagement software provide end users with both great challenges and opportunities.

The trends we see today basically have one common theme: telemanagement is still a relatively new technology, and therefore, users seeking management software are not always sure of the applications the system will have to support. A system purchased to run on a microcomputer may eventually need more capacity or processing power as the organization grows. Built-in reporting functions may not provide all the types of information users may need once they become more familiar with telemanagement opportunities.

The Telemanagement Symphony

With increased automation and centralized management comes a need to pass data to other systems or to consolidate billing at a central mainframe location. With the growth of the telemanagement industry and the increased capabilities of the hardware and software itself, users are seeking new strategies to implement telemanagement.

NEW LANGUAGES

Very high-level languages, also known as fourth-generation languages, such as Informix, Oracle, PRO IV, and Paradox, are machine-independent languages that can run in many different operating environments. Application development can proceed more quickly because the programmer is more concerned with the functionality of the program than with the technical aspects of the hardware on which it runs.

To the end user, an application written in a fourth-generation language provides easy access to data, more control to set up applications, and extensive query functions.

The use of fourth-generation languages provides maximum flexibility by allowing ad hoc report generation and ad hoc inquiry. In telemanagement applications such as inventory, work order, cable, and trouble tracking, the relational query functions can be very desirable. Fourth-generation languages also facilitate portability, the ability to change operating environments with a minimum of disruption.

While the fourth-generation language is quickly becoming an industry buzzword, users should remember that certain functions, such as call costing and sorting, can be better addressed with traditional hierarchical data base structures. Certain vendors have begun to mix applications in different programming languages to better facilitate the functions.

PORTABILITY

Today's rapidly changing technological environment and increasing reliance on data obtained from application programs make portability increasingly important. Portability can refer to either a portable language (such as a fourth-generation language) or a portable software package system that can run on different types of computer hardware. When looking into a product that is said to be portable across environments, users must ascertain whether the system is written in a portable language and whether the vendor has truly tested and documented the system in each of the environments.

Portability has obvious advantages; it allows the system to grow with the organization. But potential buyers must assess the integrity of the product to ensure that the system and support do not try to be all things to all users.

DISTRIBUTING THE DATA

End-user market studies conducted by The Lido Organization, Inc. reflect a definite increase in the use of distributed data processing in many organizations. For example, while call-accounting functions commonly reside on a mainframe, running facilities management functions at the mainframe may not be desirable. The processing functions can be made easier either by using a system interface that can pass data to the mainframe without redundant entry or by distributing the functions of telemanagement across various operating platforms.

Additional studies indicate growing requirements for interfaces into other information systems such as links to personnel, general ledger, accounts payable and accounts receivable, security, and network management.

The emergence of electronic data interchange services and standards will also greatly affect the telemanagement industry. While a good deal of work has already been proposed regarding electronic billing by vendors, end users see benefits in electronic transfer of service orders and trouble tickets. This capability could speed resolution as well as keep the client informed of delivery dates and escalation procedures.

In a technological environment that is quickly becoming more and more computer-driven, the ability to distribute data across many domains gives the end user even greater flexibility and enhances the level of management control over critical communications systems and their vendors.

In the management of corporate information networks, growth in decentralized operations accompanies growth of centralized control and management. The need for both environments also stimulates the need for distributed data processing as applied to telecommunications and network management.

CONCLUSION

The number of decisions today's telecommunications manager must make is increasing every day. Technology is expanding the corporate role of telecommunications, giving the telecommunications manager more of a say in setting strategic directions for the

The Telemangement Symphony

company. But the merging of computers and communications makes it necessary that the telecommunications manager acquire yet another level of knowledge—the nature of data processing approaches and computer technologies.

Both telecommunications professionals and traditional MIS staff must embrace the new technologies and their possible uses. If a proper decision regarding telemangement system configuration is made from the outset, that system will not only enhance the telecommunications function, it will also advance the career of the individual implementing it. □

Management of Centrex Systems

This report will help you to:

- Evaluate customer-controlled Centrex network management capabilities.
 - Identify three major management applications required by Centrex users.
-
-

MANAGEMENT PACKAGES AND SYSTEMS

There is clearly a trend with Centrex services to provide a growing array of administrative and management tools to the customer. This trend is part of an attempt by the telcos to match, with their Centrex service offerings, the features and facilities that come with digital PBX systems

The management tools are presently available from two sources: the telco rents or sells services and systems to its Centrex customers; and third parties sell systems that can be added to a Centrex installation.

We can identify three major management applications needed with Centrex: *system administration*, which looks after additions, moves, and changes (any may also include the attendants' consoles); *system management*, which is primarily concerned with call detail recording, cost allocation, and optimization; and *network management*, which is associated with problem determination and network configuration.

These three significant Centrex management requirements are addressed in more detail, with some examples, in the following sections.

SYSTEM ADMINISTRATION

A prerequisite to a satisfactory arrangement for customer-controlled administration of a Centrex system is to have a structured, customer-owned, telecommunication cabling system. This cabling system should provide jacks with a sufficient number of twisted wire pairs wherever a voice or data terminal is ever likely to be placed. The objective of such an installation is to plan for a system lifetime of ten years and to be able to handle all of the physical additions, moves, and changes without extensive re-cabling. This type of approach requires a fairly high investment in a well-planned cabling scheme, but can save several times its initial cost by minimizing new cable installation over the ensuing years. Another part of a well-planned cabling installation is the use of lockable wiring closets, provided on the basis of at least one per floor, to house all cable terminations and cross-connect arrangements together with any electronic hardware, such as controllers, that needs to be provided on a floor.

Several major suppliers now offer universal cabling systems that are based primarily on the use of twisted-pair copper wiring, avoiding the use of any coaxial or twin-axial cable, and that employ optical fiber wherever that may be justified. One of the best in-building cabling systems is the Premises Distribution Systems (PDS) from AT&T. The PDS includes patch panels and plug-ended cords for cross-connecting twisted-pair wiring and optical fiber ca-

This Datapro report is based on Chapter 3, "Management of CENTREX Systems," from *The Manager's Guide to Centrex*, by John B. Abrahams. © 1988 Artech House, Inc. Used by permission.

Management of Centrex Systems

bles. The AT&T 110 Patch Panel system provides for cross-connections of multiple circuits, in 300 and 900 pair cable sizes.

Northern-Telecom offers its Integrated Building Distribution Network, which is a well-planned and fully documented system, also based on copper wire pairs and optical fiber.

A third alternative is the structured cabling system that was designed, and is still recommended, by IBM, now available from independent suppliers. The IBM cabling system is unique in its emphasis on shielded twisted-pair wiring, which can be valuable to ensure higher bit rates or lower error rates. Unfortunately the IBM-designed cabling system employs bulky plugs and sockets for its cross-connect panels and is usually three times as expensive to install as most competing systems.

A variety of other, similar, universal voice-data cabling systems is available from the telephone operating companies and from independent contractors. Regardless of the source of the wiring, this matter must be emphasized much more than it was in the days of simple analog Centrex systems. It deserves careful design, planning, and installation, together with complete electrical testing of each of the circuits after installation and before being cut over for regular use.

The AT&T 5ESS offers on-site operations, administration, and maintenance (OAM) capabilities through an on-line terminal. Change and verification of the data base in the 5ESS may be done from the customer's offices, at the central office, or from a remote site. When this responsibility is shared by a system administrator employed by the customer and by telco personnel, there must be very close cooperation between the two groups. If the two OAM groups are not working in a coordinated way then it is better for the customer not to be involved in system administration at all.

Northern Telecom has announced its Customer Site Administration (CSA) package, although this has not been implemented as yet with many digital Centrex customers. CSA is a computer-based subsystem that provides the customer with the ability to rearrange telephone sets and features. It requires a personal computer in the system administrator's office to interact with three modules in the DMS-100 at the central office, namely the Master Control Unit, Gateway, and Customer Assistance Center. The CSA package includes a system audit capability to keep a record of all adds, moves, and changes that are made in the data base, as well as a set of self-diagnostic reports.

The EWSD ISDN Centrex system from Siemens will include a customer station rearrangement (CSR) capability that will be run from an administrative terminal on the customer's premises. CSR will be used to control the move of stations to different locations and to change the allocation of Centrex features as needed. Various management reports will be available from CSR to provide the administrator with the necessary information to control the Centrex ISDN environment.

Several independent companies now offer Centrex management systems for administrative, as well as call-accounting, applications. One of the better known of these systems is Cenpac from American Telecorp, which is in use on systems with a total of over one-half million Centrex lines. Cenpac is compatible with 5ESS and DMS-100 central offices (as well as with the older 1ESS and 1A ESS). This package automates line and feature changes, gives a user control of the Centrex data base and provides a number of user-definable reports.

Other aspects of system administration that have been addressed by a few suppliers are the attendant's console and message center. For example, Conveyant Systems has announced the Teledesk Centrex Workstation, based on a PC-AT. This unit provides enhanced call processing by substantially reducing the number of required keystrokes; an electronic directory; and a message center.

Most telcos make an extra charge to users who employ a system administration terminal on their own premises. New England Telephone (part of NYNEX) has announced that with its Intellipath II (digital Centrex) service in Massachusetts there will be no cost to use its Customer Line Administration feature.

SYSTEM MANAGEMENT

System management is usually associated with the major function of call accounting, which involves the collection of station message detail records from the Centrex system and the processing of SMDR to produce call detail records. The CDR output can be distributed to departments or to individuals, to highlight telecommunication expenditures and to provide for charging back, if that is corporate policy.

Until very recently the SMDR and CDR capabilities that were available with Centrex services lagged seriously behind similar systems and services that are widely used with digital PBX systems. This situation is now starting to improve, but there is still a good

Management of Centrex Systems

deal of dissatisfaction among telecom managers with the CDR support from some telephone operating companies.

The common approach to call accounting with digital Centrex is for the telco to supply SMDR on magnetic tape from the central office system, on a monthly batch basis to the customer. The Centrex customer, in turn, then sends the tape to be processed by a telecommunication service bureau, which prints monthly reports for distribution within the customer's organization. Although this procedure is better than no call accounting at all (which was essentially the situation with older Centrex systems), it still falls far short of the needs of a well-managed, telecom-conscious organization.

A call-accounting subsystem should enable the Centrex system administrator to make on-line inquiries (e.g., CDR for a 24-hour period or for a specific department); to define custom-tailored reports; to set the start and finish dates of a reporting period to coincide with the telco's actual long-distance billing dates; and to obtain a wide range of traffic data. In this way the SMDR output can be used to optimize the user's network, in such areas as WATS, 800 lines, and leased circuits, as well as to perform cost accounting.

The 5ESS has a separate applications processor (AP) that provides message detail recording to customers' locations. The MDR data are transmitted in real time from the main 5ESS processor to the AP over a packet-switched link. The AP stores these records and can process them to produce reports on all originating calls going over any type of private outgoing trunk, such as ETN, CCSA, WATS, or FX. Incoming call reports give details of all tandem-type calls coming in from other nodes on the customer's network. Public MDR reports are produced for all calls from the Centrex system into the public switched network. Since the AP is on the telco's premises, it is unlikely that the user has any control over the format or content of call accounting reports produced by the 5ESS.

Northern Telecom has not yet delivered any improvement to the features of its Customer System Management, which provides heavily preprocessed data tapes from the DMS-100 system. The Dynamic Network Controller has been announced by Northern Telecom as a system to deliver on-line call detail records to customers' premises but this is not yet on general commercial release.

A few independent suppliers are taking advantage of the tardiness of the major Centrex manufacturers and

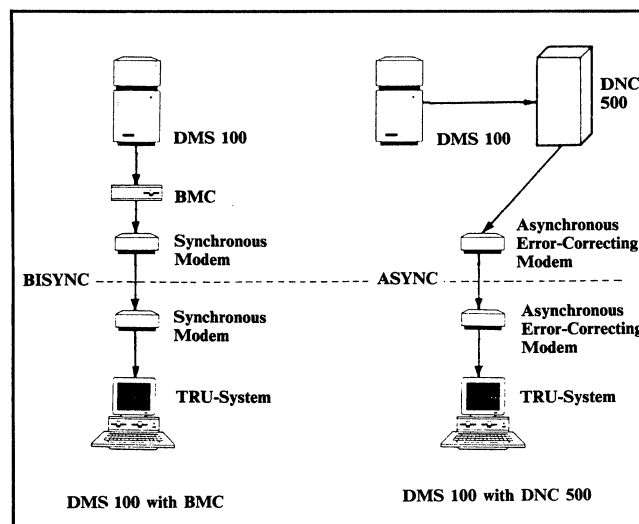


Figure 1. TRU System with Meridian Digital CENTREX.

are selling powerful CDR systems to be attached to a Centrex system, whether based on a 5ESS or DMS-100.

Telco Research has been selling its TRU system for CDR for over three years to Centrex users who are on analog No. 1 and 1A ESS. It now has several options for digital Centrex users. Customers served by Meridian Digital Centrex may use either a synchronous or asynchronous data link between the DMS-100 in the central office and an IBM PC-AT or PS/2 at the user's location. These alternatives are both illustrated in Figure 1.

The synchronous option requires a Billing Media Controller (BMC), which is supplied by the Cook Electric subsidiary of Northern Telecom. The BMC emulates a tape drive and uses bisynchronous protocol over a data link, at 2.4 or 4.8 kb/s, for polling from the TRU system in the system administrator's office.

An alternative is to use the DNC, which is a new Northern Telecom system that can supply real-time SMDR or CDR data to multiple Centrex customers. In this case, asynchronous data communication is used, up to 9.6 kb/s, and call detail records are stored and processed in the TRU system.

AT&T's 5ESS Centrex system provides real-time SMDR data from an AT&T 3B minicomputer. With this system an asynchronous link with error-correcting modems is also required. The 3B processor has the advantage that it can buffer data in the case of transmission failure. When the data link is restored then the 3B resumes the transmission of call detail to the TRU system at the customer's premises. This configuration is shown in Figure 2.

Management of Centrex Systems

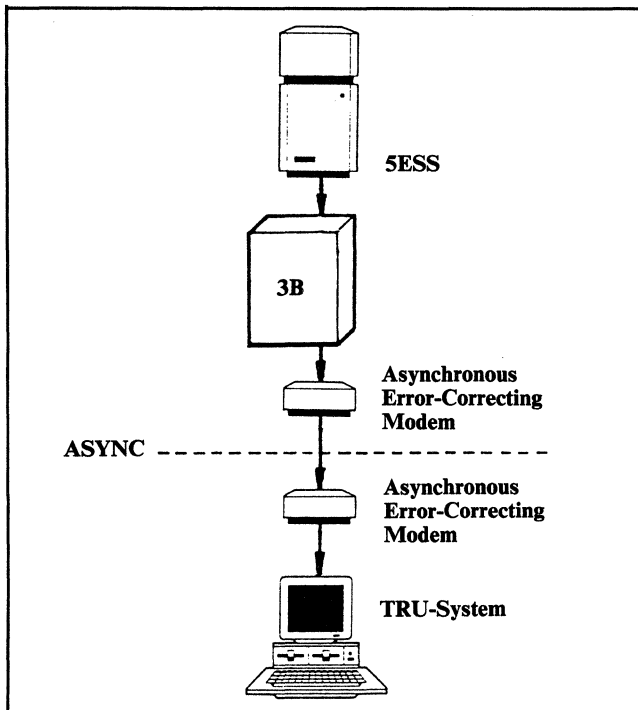


Figure 2. No. 5ESS with 3B computer for SMDR and CDR.

Another example is the M3000 Centrex Telephone Cost Management System (TCMS), which is supplied by Moscom Corporation. So far this software has been used to process CDR from the AT&T 5ESS. The TCMS program and data base resides in a personal computer at the customer's site and communicates with the 3B2 computer at the Centrex serving office over a dial-up line. TCMS also incorporates an automated directory and a message center capability. The cost of this software ranges from two to six thousand dollars, depending on the number of Centrex lines that are supported. Call records may be retrieved on an on-line basis, as well as both summarized and detailed call-accounting records.

NETWORK MANAGEMENT

Network management is concerned with the minute-by-minute operation of a Centrex system or network; the collection and classification of error messages (diagnostics); the testing of circuits and the restoration of service after a breakdown; and the collection of traffic data in order to optimize the network and plan for

the future. For these applications, as with system management and system administration, the trend is very much toward putting these tools directly into the user's hands, although the great majority of Centrex customers have not yet moved very far toward in-house network management.

AT&T has designed and implemented complete Switching Control Center Systems for a few large Centrex users.

Northern Telecom has announced the DNC-100 (known as the NM-1 in the SL-1 PBX environment) for network management from the customer's site. This system provides real-time network monitoring, with problem isolation, the management of problem resolution, and performance reporting. Network analysis and network design programs are becoming available for use with this system. The DNC-100 is intended to produce automated trouble tickets, keep a trouble ticket history file, and provide for customer-defined exception reporting.

Network Management is now recognized as a vital aspect of any major telecommunication system, but the provision of this capability to the users of Centrex is still in its infancy.

AT&T has announced the concepts of its Unified Network Management Architecture (UNMA), but has yet to deliver any software to implement this plan. *[EDITOR'S NOTE: AT&T's UNMA protocols, as implemented in ACCUMASTER network management products, were introduced in early 1989.]* An increasing number of high-volume users are taking advantage of software defined networks (SDN), in which the availability of such services as WATS and 800 lines can be adjusted dynamically to accommodate varying traffic loads. The SDN can be implemented in conjunction with Centrex.

It is likely that the Netview architecture, designed and being implemented by IBM, will be the dominant approach to integrated voice-data network management. The integration of Netview with central office systems and so with Centrex service is likely to start soon, probably with switches from Siemens or Ericsson, because IBM has signed cooperative agreements with these two major telecommunication system manufacturers. □

The Evolution of Automated Management Systems in Voice Networks

This report will help you to:

- Trace the progress of voice network management and prepare for future developments.
 - Select the network management system best suited to your organization's needs.
 - Make the transition from manual to automated network management systems.
-

The purpose of automated systems is to minimize human involvement in relatively mundane or redundant tasks, thereby reducing the chance of error due to inattentiveness, carelessness, and deviations from the norm. Automation also speeds up and adds precision to complex tasks, possibly eliminating human involvement altogether. As applied to network management, the degree to which various operations can be automated has a direct bearing on the performance and cost of installing, operating, and maintaining information systems and communication networks.

THE EVOLUTION OF AUTOMATED SWITCHING SYSTEMS

Early Automation

Voice network management techniques began with the transition from manual to automated systems. In 1878 the first operator-controlled switchboard was implemented by the Bell Telephone Company to serve 21 subscribers in New Haven, Connecticut.

This report was developed exclusively for Datapro by Nathan J. Muller. A former consultant, Mr. Muller has 18 years' experience in the computer and telecommunications industries. He has written extensively on all aspects of computers and communications and is the author of *Minimum Risk Strategy for Acquiring Communications Equipment and Services* (Artech House, 1989).

The operator had full responsibility for answering call requests, establishing appropriate connections, and breaking them when the calls were completed. The operator interconnected subscribers through manual cable connections at a patch panel.

As the number of telephone subscribers grew, so did the physical size of the patch panels. By the 1880s, hundreds of operators were required to serve customers in major metropolitan areas. Entire floors of multistory office buildings accommodated row after row of these devices. Managers actually wore roller skates to supervise operations. In 1889, the Strowger switch, named after its inventor, Armond Strowger alleviated some of the unwieldy requirements of telephone switching. The step-by-step switching device eliminated the requirement for so many operators by automating local call handling.

Such levels of automation were eventually applied to the PBX, which was still essentially an operator-controlled patch panel reminiscent of the first central office switchboards. In 1965, telephone companies started replacing electromechanical switches with electronic switching systems that operated under stored program control. This type of technology, too, was eventually applied to the PBX.

The Evolution of Automated Management Systems in Voice Networks

The Automated PBX

Today's PBXs combine stored program control, advanced processing power, large-scale circuit integration, and high-capacity memory to support an incredible number of features. Although more compact and sophisticated in design, they provide the same basic functionality as the first generation of switchboards. The difference is in the process of receiving call requests, setting up the connections, and tearing down the paths upon call completion—they are entirely automated, which means that more calls can be handled in less time.

As the volume of outbound and inbound calls increased, other manual tasks became the targets of automation and were integrated into the PBX. Speed dialing, for example, allows the user to complete calls by dialing an abbreviated number. Automatic call distribution allows sharing of incoming calls among a number of stations so that the calls can be served in the order of their arrival. Automated attendant is the capability of the system to answer incoming calls and prompt the caller to dial the appropriate extension or leave a voice message without using the operator.

Automated PBX Management

Call Accounting

Automated PBX call management systems were introduced in 1972 by Account-A-Call (Los Angeles, California). Commonly referred to as "call accounting" systems, these products consist of specialized hardware and software that collect call records from the PBX, translate and organize the data, and generate printable management reports. Such systems were of particular value to firms in the service industry that needed to identify calls made to and on behalf of clients for billing purposes and to identify all calls made by various departments for allocating costs.

Until call accounting systems appeared, a common method for processing telephone bills involved assigning a secretary or accounting clerk to review the company telephone bill to sort calls according to the last four digits of clients' telephone numbers. That way, the company would have the documentation necessary to bill clients for the purpose of cost recovery.

As late as 1983, International Tariff Services, a Washington, DC firm serving the maritime shipping industry, was spending 61 hours of staff time per month trying to identify client calls. This manual method of

call accounting was not only inconsistent from one person to the next but resulted in too many mistakes. Calls that could not be identified by client were charged to overhead. There was no way to identify calls by department. As a result, a sizable percentage of the firm's communications costs could not be recovered, let alone allocated, controlled, or budgeted.

ITS then moved to an automated call accounting system. Four-digit account codes identified clients, projects, or the appropriate department making the call. Previously, these account codes had to be entered by each employee, or the call would not go through. The call accounting system provided monthly call detail reports, which gave a complete breakdown of calls by client, department, and project—all sorted by account code.

With this information at hand, it was a simple matter to incorporate call costs into each client's invoice for services. With all calls fully accounted and billed for, very few had to be absorbed as overhead. The number of unauthorized personal calls also dropped dramatically, saving the firm even more money. In moving from a manual to an automated system, ITS staff spent only 1 hour and 15 minutes per month in telephone bill processing rather than 61 hours—a savings of 59¾ hours. This increased the efficiency of staff and eliminated the need for overtime.

Most call accounting systems today report trunk usage. This information can be used to configure a more efficient network. For example, such reports can be used to determine the need for additional tie lines to high-traffic locations, document the need for higher band WATS lines to extend coverage, or justify the need for high-bandwidth, bulk-billed digital services like AT&T's MEGACOM, MCI's PRISM, and US Sprint's ULTRA WATS. By noting low usage, these types of reports can also be used to identify unneeded trunks or possible trunk outages at remote PBX locations.

PBX System Management

Managing PBXs, however, is more than just call management. Since their introduction, call accounting systems have been enhanced with features and functionality that extend far beyond mere call detail reporting to include the automation of maintenance tracking, inventory and accounting, facilities management, moves and changes, circuit usage monitoring, trouble reporting, and much more. Many of these innovations were spurred by the breakup of AT&T in 1984. With the fragmentation of the telecommunications industry, users had to take more responsibility for managing their communications resources. Sens-

The Evolution of Automated Management Systems in Voice Networks

ing new opportunities, CPE vendors went about automating just about every task traditionally provided by "Ma Bell."

The PBX system administrator's job is very much like that of a computer center operator's: management, control, and diagnosis of the PBX and its network of attached devices. This includes the following:

- Making changes to the system database to add or update station information, users, and classes of service.
- Requesting statistics and status reports.
- Ordering diagnostic tests and analyzing the results.

Permission classes and multilevel passwords are typically used to limit administrative access. Operating a PBX is similar to operating a computer, the difference being that the PBX' application is narrow in scope. The application of maintenance and programming (MAP) output channels, however, has steadily evolved toward more comprehensive applications. A circuit assurance system (JANUS) interfaces to a PBX MAP port, for example, on which the PBX can be used to monitor the performance of all lines in the system, providing comprehensive line outage and trouble reports. A number of systems use MAP channels to support more comprehensive ACD traffic reporting in both realtime and historically to develop trend information. The MAP channel can also provide an interface to more sophisticated network management systems, which can control an entire voice/data network. Network management systems—such as Cincom's NetMaster, IBM's NetView, and Dynatech's Prism—serve as information collectors that act on alarm conditions and monitor the progress of the user's entire network. In this regard, there has been a more proactive application of PBX MAP channels in recent years than was true in the past.

MANAGING CENTREX

After AT&T divested the Bell Operating Companies (BOCs), it was up to them to revitalize their Centrex offerings. After adding more features, the telephone companies began to address the area most criticized by prospective customers—lack of management and control.

In addition to obtaining call detail on a virtual realtime basis (instead of waiting for billing tapes), users can now control which numbers, features, services, and billing codes are assigned to each line via an on-premises terminal and interactive software pro-

gram. Instead of waiting for changes to be implemented in a matter of days or weeks, they can be implemented in a matter of minutes. Even adding or deleting lines falls under the automated management capabilities of Centrex.

Not only can the network manager review the status of the current Centrex configuration, but plan ahead to meet future demands by determining the effective dates of the changes. All such configuration information is stored at a management system located at a central location in the telephone company's serving area. Individual Centrex exchanges poll the system for any customer changes, which are then loaded in the master database. This causes internal telephone company records to be automatically updated.

MANAGING T1 MULTIPLEXERS

There are a number of features available with T1 multiplexers that automate bandwidth sizing, circuit configuration, and disaster recovery. For example, when circuits are added to the network, they are automatically routed when the operator enters the end points. The circuits to be routed are identified by the operator on a per circuit or user group basis. The best path is chosen based on the match between the circuit profile and on the attributes and parameters of the aggregate.

The characteristics of each aggregate are determined during network configuration. The data entered is used in conjunction with circuit profiles to insure optimum routing. In this way, a route is performed based on quality considerations to insure the integrity of the applications. The quality-based parameters of each aggregate include delay, error rate, availability, and user-defined attributes.

A profile is created for each circuit, allowing the network manager to set options for assigning priority for bandwidth, optional manual routing, and whether downspeeding will be allowed or not allowed during line failures. A set of qualifiers may be established, which are used to automatically route circuits over the correct aggregate types. These qualifiers include mandatory, desirable, undesirable, not allowed, and "do not care." The shortest path between the two circuit end points that meets the requirements specified in the circuit profiles is used.

In addition to automating the routing function, circuits may be downspeeded during rerouting to ensure that all users continue to communicate, rather than just a few. In offering software-selectable optioning of ADPCM at 32K bps, 24K bps, and 16K bps, adaptive downspeeding can be implemented to keep users on-

The Evolution of Automated Management Systems in Voice Networks

line during intelligent automatic rerouting instead of letting them get bumped off. This enables continued operation with efficient T1 bandwidth fills.

The system automatically calculates rerouting based on each likely failure. In the event of a failure, the system automatically recalculates optimized routing based on current network conditions. After restoration, the system again automatically calculates the best rerouting should a second failure occur on the network. The system will then be ready to handle the next emergency until there is not enough of the network remaining through which to reroute traffic.

The ability to reroute a full T1 circuit without timing out front-end sessions is critical. Rerouting a single circuit can be done in as little as 200 ms. For 25 circuits, this works out to only 5 seconds, well within the time-out thresholds of front-end sessions. This ensures that users do not have to manually restart front-end sessions after data circuits are rerouted.

High-capacity digital networks must also adapt to accommodate changing application needs. Time-oriented reconfiguration allows users to do so on a scheduled basis. Circuit routing may be altered to accommodate applications that change from day to night. Since voice traffic tends to diminish after normal business hours and data traffic changes from transaction-based to batch, the management system provides the means to adapt automatically, without operator intervention.

In other cases, organizations might want to alter their networks to track the business day around the world. Circuit end points can be adjusted automatically in order entry applications, for example, allowing new order entry terminals on the West Coast come on-line while those on the East Coast shut down. That way, business is not lost; all calls get answered.

MANAGING CROSS-CONNECT SYSTEMS

Within the context of customer premises equipment (CPE), a cross-connect system is a front-end processor to a tandem PBX. Because the cross-connect system provides the means to "nail up" connections, it permits the more efficient usage of PBXs for other memory-intensive call handling features, as well as sophisticated options such as voice mail, automatic call distribution, and automated attendant.

The basic function of the cross-connect is to accept aggregates of channels via T1 facilities and groom them for individual routing. Since connections are defined in software, reconfigurations may be imple-

mented in a matter of minutes from an authorized terminal, thus automating the entire process of circuit provisioning.

An alternative to CPE is to subscribe to the Customer Controlled Reconfiguration (CCR) service of local or interexchange carriers. This permits telecom personnel at the company's main office the ability to alter the circuit configurations as needed to control the cost of telecommunications. During peak hours, for example, the cross-connect may be configured at a CCR terminal to make more circuits available to accommodate high usage. If traffic reports indicate trends in usage by hour-of-day or day-of-week, this information can be stored in a microcomputer. Primary, secondary, and tertiary configurations may then be uploaded to the cross-connect system with only a few keystrokes.

EXPERT SYSTEMS

With automation slowly but surely diminishing the need for human intervention, some vendors have sought ways of applying expert, or knowledge-based, systems to network management, particularly to managing multivendor networks. Developments in resolving this problem could result in much faster diagnostics, less network downtime, and considerable cost savings in technical personnel.

Expert systems apply rules to a given set of adverse circumstances to arrive at the best possible solution, thus simulating elementary human decision-making. The knowledge base and the rules for effectively using that knowledge are, of course, supplied by human experts. Consequently, the capabilities of the expert system are limited by the knowledge base it can draw upon and the rules governing its use. Although such systems can greatly enhance the user-friendliness of network management systems, they are not yet capable of learning by themselves. For this breakthrough to happen, considerably more research must be done on the fundamental questions of what it means to perceive, reason, and learn in the first place.

Nevertheless, expert systems hold immediate promise in such areas as configuration management, disaster recovery, and trouble ticket processing. In fact, any routine task is a likely target for this level of automation.

Another area in which expert systems can make a difference in network management is status reporting. While vendors pride themselves on their equipment's abilities to provide comprehensive status information on both equipment and links, network managers have become lost in a blizzard of paper. There is now more

The Evolution of Automated Management Systems in Voice Networks

information available about various network operations than can be effectively assimilated or used. Much of this information is redundant, since the equipment or link status may not change significantly, if at all, on a per-minute basis, or even from hour to hour. Expert systems can not only remove redundant information before it is output, but improve the analysis and presentation of data in the form of meaningful, relevant management reports.

As more and more vendors adopt IBM's NetView for SNA-type networks and AT&T's Unified Network Management Architecture (UNMA) for end-to-end management over wide-area networks, users will be able to better organize their networks. As the various network components become better integrated, many more possibilities for network management and control by expert systems will open up.

MAKING THE TRANSITION

Today's network management systems have demonstrated their value in permitting technicians to control individual segments or the entire network remotely. In automating various capabilities, network management systems can speed the process of diagnosing and resolving problems with equipment and lines. Combined, the capabilities of network management systems permit maximum network availability, thus enhancing the management of geographically dispersed operations, while minimizing revenue losses from missed business opportunities.

Despite these advantages, it is still difficult for many managers to appreciate the time, effort, and expense involved in making the transition from manual to automated systems. Even harder to grasp is the level of ongoing commitment that is usually necessary to maintain such systems.

In the realm of call accounting, for example, tariff tables must be installed and continually updated to keep pace with changing rates and new services. A database describing the configuration of the telephone system at each node of the network must be set up and continually managed. The accuracy of input data and report runs must be verified; all time-consuming tasks requiring a high degree of expertise and staff continuity.

When multiplexers are used to integrate and manage voice and data over the wide-area network, automated management and control capabilities can preempt adverse situations to insure maximum network availability. For this to happen, however, a database must be created consisting of a complete profile for each circuit. The profile includes such items as the interface definition (e.g., asynchronous, synchronous, bisynchronous, etc.) and the data rate. Various options for assigning priority for bandwidth, optional manual routing, and whether downspeeding will be allowed or not allowed during line failure scenarios must also be entered into the circuit profile. Since the only thing about networks that stays the same is change, it is unavoidable that a certain amount of staff time will be devoted to database maintenance.

Even knowledge-based management systems require a high degree of involvement at the start of implementation and an on-going database maintenance routine. After all, the knowledge base must be compiled and the rules written in machine-readable format so that the inference engine can make valid decisions and take appropriate restoral actions.

Although relatively inexpensive and reasonably powerful expert system "shells" are available to help in assembling knowledge-based systems without delving too deeply into the intricacies of LISP or other artificial intelligence languages, when new situations arise, the knowledge base must be expanded and new rules added. Not only does this take time, but it takes a level of staff expertise not normally found in today's network control centers. The real value of expert systems, as applied to network management, may not be in eliminating personnel but in their ability to eliminate extraneous information to arrive quickly at the crux of the problem and, once defined, recommend or implement a solution.

Even though "automation" implies that the organization can improve performance at considerably less cost in manpower, it also introduces a new level of complexity, usually requiring that present staff augment their present levels of expertise with additional training. Depending on the network management system, automation may even require the addition of highly specialized personnel. Consequently, to automate network management merely to trim staff may be too simplistic an objective. Any staff reductions that result are most likely to occur over time as a by-product of automation. □

The Missing Link— Network Management

This report will help you to:

- Compare features common to most digital switches on the market today.
 - Assess the effectiveness of network management controls for network stress conditions.
 - Use network management systems to satisfy changing demands on DOD communications systems.
-
-

One of the major problems Military Departments (MILDEPs) and other Department of Defense (DOD) entities are facing in today's environment is the rapidly changing and dynamic environment of existing near-term and future user demands on their communication systems. A key element in satisfying these demands is network management, which has lagged behind the technological advances of switching and transmission equipment subsystems.

INTRODUCTION

The DOD is undergoing dramatic changes by taking advantage of new developments in communications technology. The switched voice subsystem is being expanded and modernized, through the use of commercial digital technology, to replace the existing analog Automatic Voice Network (AUTOVON) and other local and MILDEP switches. The overall goal for the DOD management support structure is to maintain network performance regardless of network traffic overloads or partial network disruption or destruction. Failing this, the goal is to assure network availability at

least to high-priority users. Network management is critical to both the peacetime and crisis/wartime accomplishments of these goals. Current control capability is limited primarily to a few manual controls, while new digital switch technology provides both automatic and remote control capability as well as an increased number of controls. Further, the remote capability of these switches will enhance service restoral and reconstitution; centralized Administration, Operations, and Maintenance (AO&M); and real-time network monitoring. Therefore, the approach of using commercially available digital switches requires tradeoffs in network management requirements. Automatic and manual control functions, as well as the interswitch signaling scheme, have to be identified and specified to support the wartime mission of the DOD.

This report assesses the effectiveness of network management for various network stress conditions and surveys the applicability of current digital switches to support network management requirements. It is based on a study performed for the Defense Communications Engineering Center, Contract Number F19628-82-C-001.

CONTROL ASSESSMENT

The primary technique used in analyzing the effectiveness of controls was event-by-event simulation. A discrete-event computer model (provided by the De-

This Datapro report is based on "The Missing Link--Network Management," by Donald J. Jurenko and Robert L. Sligh, Jr., the Mitre Corporation, from the Military Communications Conference 1987, Washington, DC, October 19-22, 1987. © 1987, IEEE. Reprinted with permission.

The Missing Link—Network Management




Figure 1. Network configuration.

fense Communications Agency (DCA)), which provides a detailed representation of the functions of circuit switched networks, was used to simulate network conditions. As a result, the effects of proposed changes in network configuration and routing, traffic volume and distribution, and switch hardware, as well as trunk and switch outages, could be analyzed. A 27-node network, as shown in Figure 1, representing a portion of the Defense Switched Network in Germany was modeled.

The model simulates calls between individual network subscribers or on a switch-to-switch basis, switch common equipment, and the different types of routing logic used in the actual network. The simulator functions by modeling each of the discrete events that occur as calls are routed through the network. Call processing is simulated from the time the call originates at a subscriber until it terminates at the end of conversation. All significant events that occur in the real system during call processing are modeled. These include obtaining an access line, dialing, queuing for common equipment at the switching center, alternate routing, seizing trunk connections, preemption, inter-switch signaling, call blocking, reattempts, and conversation time.

Network stress conditions evaluated were (1) general traffic overload, (2) a change in traffic characteristics, (3) focused traffic overloads, (4) switch outage, and (5) trunk outage. Both common channel signaling (CCS) and inband signaling were used in evaluating the general traffic overload condition; only CCS was used for the other stress conditions. Emphasis was placed on signaling because of the long transition time expected for phasing out inband signaling.

Traffic flow controls analyzed were access control, route control, trunk reservation, and code cancellation. The methodology used in assessing control effectiveness was to:

- Evaluate the performance of a network under normal operating conditions (baseline)
- Apply stress conditions to the network
- Analyze the resulting network performance and apply controls to improve user service

Significant results of the study were:

- For all control applications analyzed, the performance of the network was better when CCS was used for interswitch signaling.
- When CCS was used for interswitch signaling, network degradation was more graceful under stress conditions. More time is provided for controller analysis, issuing commands to invoke controls, and adjusting thresholds. When inband signaling was used, alternate route control and traffic throttling had to be applied immediately to improve network throughput.
- For all control applications analyzed, alternate route control degraded network performance when CCS was used for interswitch signaling. However, the opposite was true when inband signaling was used.

The results of control applications for a general traffic overload when CCS is used for in-switch signaling are presented in Figure 2. For a 25 percent increase in traffic with no control actions taken, the network GOS was P.17, with a source-to-destination (S/D) pair GOS range of P.00 to P.70. When controls were applied, the network GOS improved to P.06, with an S/D pair GOS range of P.00 to P.40. The perfor-

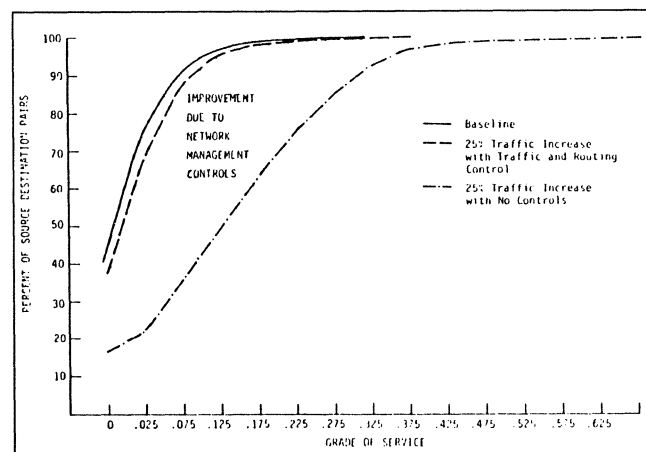


Figure 2. Control effectiveness for a general overload condition.

The Missing Link—Network Management

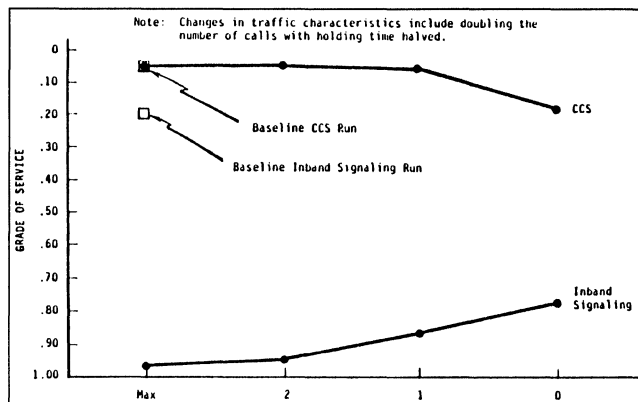


Figure 3. Network GOS for a change in traffic characteristics.

mance of the network with no stress conditions is presented in this figure for comparison. The improvement in service was accomplished without blockage of high-precedence calls, significant degradation in calls carried, and locking users totally out of the network. The controls applied were:

- Global Route Control—this can be applied automatically by switches in the network, and is based on analysis and processing of data for local and inter-switch traffic and congestion.
- Call Gapping—this can be applied automatically by switches and is based on an analysis similar to that performed in global route control.
- Network Routing Modifications—these are changes to the network routing scheme that require a control facility with total view of the network.

Similar results were obtained when the network was subjected to a change in traffic characteristics, focused overload, switch outage, and trunk outage. In all cases where CCS was used, the applications of controls improved user service while providing nonblocking service for high-precedence call attempts.

Figure 3 shows the results of the change in traffic characteristics analysis. There was no significant change in the network GOS for the normal traffic load (baseline) or for double the number of calls when CCS was used—P.037 versus P.035. For inband signaling, the GOS degraded severely from P.16 to P.97 when compared with the baseline run. As the number of alternate routes was decreased for inband signaling, the GOS improved and the number of calls completed increased. However, the most significant degradation in CCS network performance occurred when there was no alternate route available to complete calls.

Focused traffic overload simulations were performed on the baseline configuration using CCS. An amount of traffic equivalent to 10 percent of the total baseline

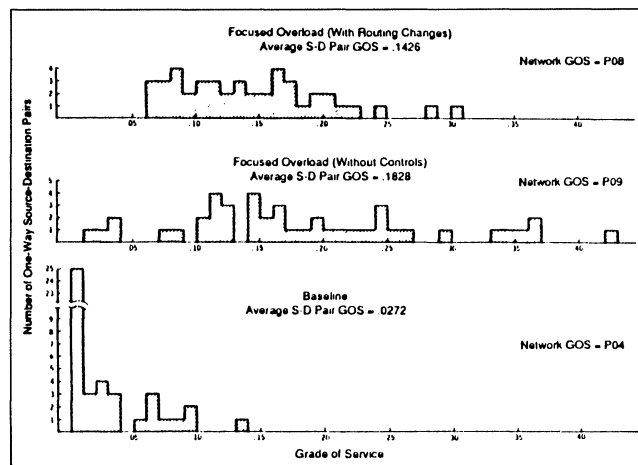


Figure 4. Focused overload—traffic to and from.

loading was added in such a way that it all went into or came out of the area around two switches in the network. The composite effect on traffic between all the source-destination switch pairs associated with one of the switches is shown in Figure 4. The performance of the baseline network without the additional traffic is provided for comparison.

The most effective controls when a switch outage occurred in the network were code cancellation and route control. The results of these control applications are depicted in Figure 5. The net effect of applying code cancellation and route control was an improvement of P.05 (P.24 to P.19). Code cancellation was applied at the originating offices for all calls destined to subscribers homed off the failed switch. The route control applied was a change in the network routing scheme. All tandem traffic or traffic that was routed through the switch destined for other switches in the network was rerouted.

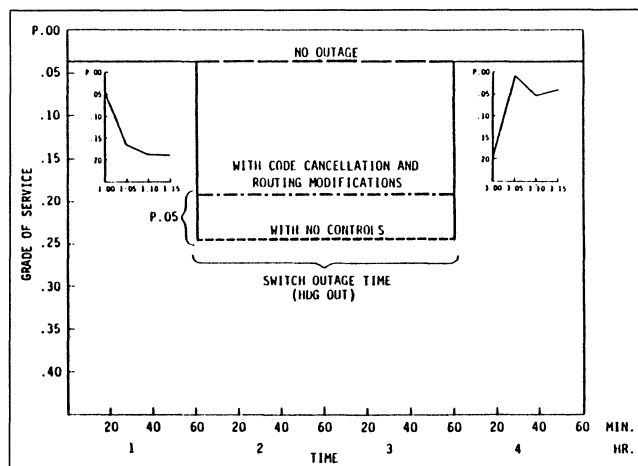


Figure 5. Network performance with a 2-hour switch outage.

The Missing Link—Network Management




Figure 6. Network performance for a trunk outage.

The analysis of control effectiveness to counteract a trunk outage was performed by removing interswitch connectivity between two switches for a two-hour period. During this outage period, the network GOS degraded from P.04 (baseline GOS) to P.11. There was also a degradation in S/D pair GOS. The maximum DOS for an S/D pair in the baseline network was P.30. When the trunk outage occurred, 41 source-to-destination pairs had a GOS greater than P.30. When controls were applied in addition to route control, there was an improvement in the network GOS and the S/D pair performance; the network GOS went from P.11 to P.07, and only 29 S/D pairs received a GOS greater than P.30. Figure 6 depicts the improvement in S/D pair performance.

Based on the results, it is recommended that:

- CCS be used in future circuit switched networks.
- The effects of the same controls applied in networks using CCS and inband signaling should differ significantly.
- At a minimum, access control (line load control, line directionalization, and call gapping), code cancellation, trunk reservation, and route control should be incorporated into all future circuit switches.
- A general-purpose network simulator should be developed to simulate hybrid signaling schemes (CCS and inband). This is especially needed because a hybrid condition will exist in the early 1990s. The controller must be aware of what controls to invoke depending on the signaling scheme used. The model could also be used to establish network control parameters and switch thresholds for specific DOD switched network configurations, to evaluate system performance, to determine further control requirements, and to assist in training controllers.

APPLICABILITY OF CURRENT DIGITAL SWITCHES

The capability of current digital switch technology to satisfy DOD switched network crisis and wartime control requirements was an important consideration in the development of the control concept. Six switches were assessed to determine their capabilities to support crisis and wartime requirements. All the switches differed in the level of support provided for network management. Most digital switches provide the basic controls, diagnostics, and tests associated with circuit switching, including internal diagnostics and tests on all equipment units, trunk directionalization, line load control, and code cancellation. Table 1 contains a list of controls, diagnostics, and features common to most present-day switches.

Other features common to all the switches include:

- Data base management
- Interactive terminals for switch maintenance and operation
- Automatic diagnostic testing of lines and trunks
- An interface for the AT&T Centralized Automatic Reporting Trunks (CAROT) system
- Trouble ticket and service order processing

The most significant finding was that there is no uniformity in application and very little interoperational capability between different switches in applying network management controls. Other findings are summarized below:

- The methods used in performing or invoking controls common to all the switches differ.
- The automatic control capabilities are not standard.
- Only three of the switches surveyed have an interface for upper-level control elements, and the interface was different for each switch.

The switches that have an interface for upper-level control also provide hardware and software for network management.

SUMMARY

All of the systems reviewed could be upgraded to perform the network management functions necessary to support the DOD requirements. Most of these up-

The Missing Link—Network Management

	Manual (Local Only)	Automated		Remote Maintenance Terminal	Regional Center Interface	Provide Central Maintenance Center
		Local	Remote			
1. Automatic Controls	X	X	X	X	Two	Three
a. Preplanned number						
b. Out-of-chain routing						
c. Selective trunk routing						
d. Short receiver timing						
e. Short sender timing						
2. Manual Controls	X	X	X	X	Two	Three
a. Group						
b. Line line control						
c. Route control						
d. Code cancellation						
e. Trunk and directionalization						
3. Switch hardware failure reporting	X	X	X	X	Two	Three
4. First level diagnosis on trunks, lines and loops (analog), and other switch equipment	N/A	X	X	X	Two	Three
5. Automatic message accounting	X	X	X	X	X	X
6. Traffic flow statistics	N/A	X	X	X	Two	Three
7. Common equipment utilization	N/A	X	X	N/A	Two	Three
8. Second and third level diagnostics on trunks, lines, and loops, and other switch equipment	X	X	X	N/A	Two	Three
9. Automatic recovery	N/A	X	N/A	N/A	N/A	N/A
10. Semiautomatic recovery	N/A	X	X	X	Two	Three
11. Switch utilization statistics	X	X	X	X	Two	Three
12. Monitor and control of other station equipment (e.g., built-in fault alarm and control system)	X	X	X	X	Two	Three
13. Interface path and test panel	X	X	X	X	Two	Three
14. Maintain station data base	Two	Two	Two	N/A	Two	Three
15. Provide real-time data base access	X	X	X	X	Two	Three
16. Monitor and control remote switching units	X	Three	Three	X	Two	Three
17. Downline loading—tables (customer)	N/A	X	X	N/A	Two	Three
18. Downline loading—software (customer)	N/A	X	X	N/A	Two	Three

Key: X = Function performed by switches.

Table 1. Digital switch features.

grades require either software modifications or additional software development.

The implementation of the network management and O&M subsystem(s) should be based on the following strategy:

- Selected existing and near-term switches should be updated to support control requirements.
- Control requirements should be incorporated into near-term and future switches.
- A switch interface unit should be used or developed to assist switches deficient in performing network management requirements.

The results of an analysis of network management techniques in a DOD environment are summarized below:

- The minimum control requirements for Defense Switched Network switches should include access control, code cancellation, trunk reservation, and route control. The switches should also have the capability to perform the analysis necessary to make control decisions and to invoke controls.
- Because routing is critical to the robustness of the network, controls, routing scheme(s), and techniques for network management should be developed at the same time. Due to the limitations of the model used in this study, only one routing scheme was used to evaluate control effectiveness. Other routing schemes (e.g., adaptive or flood) could require additional controls, as well as techniques for controlling the flow of traffic.

The Missing Link—Network Management

- In the simulations using CCS, the switch was not simulated in detail because the manufacturer would not provide CPU timing or processing characteristics. It is recommended that this information be obtained on all present and future DSN switches. The information is critical to future network management analysis because the application of controls in a network is highly dependent on the switch saturation point.
- A simulator should be developed to handle new stored program control digital switches, CCS, and automatic controls. The model should be event-by-event and have the capability to simulate different routing and signaling schemes as well as various vendor switches.

REFERENCES

- ¹M. M. Irvine, "An Electronic Watchdog for the Network," *Bell Laboratories Record*, September 1980, pp. 267-273.
- ²B. Stoffels, "Pay Now or Pay Later," *Telephone Engineer and Management*, June 15, 1982, p. 9.
- ³D. G. Haenschke, D. A. Kettler, and E. Oberer, "Network Management and Congestion in the U.S. Telecommunications Network," *IEEE Transactions on Communications*, Vol. COM-29, No. 4, April 1981.
- ⁴T. V. Greene, D. G. Haenschke, B. M. Hornbach, and C. E. Johnson, "Network Management and Traffic Administration," *The Bell System Technical Journal*, Vol. 56, No. 7, September 1977.
- ⁵Bernas and Grieco, *A Comparison of Routing Techniques for Tactical Circuit-Switched Networks*, ICC, 1978.
- ⁶C. Limieux, "Theory of Flow Controls in Shared Networks and Its Application in the Canadian Telephone Network," *IEEE Transactions on Communications*, Vol. COM-29, No. 4, April 1981.
- ⁷R. L. Sligh, *A Description of the DATRAN Network Simulator and Its Use*, IEEE International Switching Symposium, June 1972.
- ⁸*O/S VON Simulator*, Report R4942340-1-1, Defense Communications Agency, May 1975. □

Modem/Multiplexer-Based Network Management

This report will help you to:

- Trace the evolution of modem/multiplexer-based network management.
 - Grasp network management fundamentals: performance, fault, and configuration management
 - Select the right network management system for your network's particular needs.
-

Network managers are under tremendous pressure to keep their networks up and running, while controlling costs. Consequently, it is vital that a network manager select the right network management system. Although there are more network management systems to choose from than ever before, network managers now also have more problems. After divestiture, users had to assume full responsibility for their networks. This produced new problems, such as lack of trained personnel to operate the equipment in the communications center; bigger and more complex networks; and separation of network functions from the software that facilitates network performance monitoring. Many managers also face problems in planning and implementing a system that must control many different areas and/or people.

In the 1970s, few organizations had a strategy for managing their networks. Network management was sporadic; managers solved problems as they arose. Modem vendors offered some of the earliest answers to network management with their modem network management products. The modems offered built-in diagnostics for testing analog lines. At that time, the modem vendors were the major supplier of these solutions; today they are just a part of the solution. Different types of systems now play a major role in market offerings, as network management has evolved to include managing software, lines, system utilization, and other aspects.

In the early days of managing networks, vendors looked at solutions in very narrow terms: diagnostics, testing, or control. Network management systems were based on how the vendor defined network management. Early network control systems, which were primarily proprietary, dealt mainly with monitoring and diagnosing modems and multiplexers. Some of these systems offered a wraparound device that enveloped another vendor's modem—allowing the wraparound device to monitor the modem. But wraparound boxes were not the best solution; they decreased modem efficiency. So vendors looked for other ways to monitor nonproprietary equipment. The earlier systems were also limited by users' perceptions of what they wanted from the systems. Fault isolation was the main goal, so that a vendor or carrier could be contacted to come and repair the problem.

Modem/mux management systems support point-to-point, multipoint and/or multiplexed applications via a dedicated controller. Features include continuous on-line circuit monitoring, monitoring of modem conditions, built-in alarms, analog and digital testing, adaptive-rate capabilities for fluctuating line conditions, and dial restoral for automatic service restoral when leased lines go down.

All network management systems include some mechanism for monitoring the network's components. When the network management system vendor

Modem/Multiplexer-Based Network Management

also manufactures modems, the vendor usually designs the monitoring device as a built-in modem feature, eliminating the need for separate monitoring devices. On other systems, standalone monitoring devices must be attached to modems or multiplexers at each remote site. In most network management systems, these devices can monitor only physical information on the status of the modem or multiplexer, its interface with the terminal equipment, its interface with the transmission facility, and the condition of the transmission facility.

In this report a standalone, hardware-based network management system is defined as:

- *a computer-based system,*
- *owned (or leased) and operated by the user and*
- *independent both of host (and front-end processor) applications and of outside transmission facilities, that*
- *monitors the network's components;*
- *records information on those components' status;*
- *displays that information for the operators' attention and action;*
- *maintains one or more databases of network status, configuration, inventory, and history; and*
- *generates reports for management based on information in those databases.*

The computer that drives a network management system can be an on-board microprocessor, a personal computer, or a medium-to-large, dedicated minicomputer. Some microprocessor-based systems are completely modular, sometimes to the degree that a microprocessor controls and monitors each line.

Information on the modem or multiplexer and its interfaces comes from the presence or absence of signals on various EIA interface leads. Information on the transmission facility comes from measuring various analog parameters such as signal level, noise, distortion, phase jitter, and line hits. If a given interface signal or analog characteristic falls out of specification, the system's monitors set off an alarm to notify the operator of a failure.

CONFIGURATIONS

A minimal hardware-based network management system consists of a central processing unit, a hard disk

or diskette storage device, an operator's console, and a set of local and remote monitoring devices. The central processor may be a single minicomputer or may contain a number of function-specific microcomputers. The operator's console is often a color CRT, although some systems use only monochrome displays. In most systems, the operator station includes a printer; some may also include a color plotter for graphic information.

The nature of the monitoring devices depends on the native market of the system's vendor. Systems from modem vendors, such as Memotec, Racal-Milgo, Paradyne, and AT&T, use diagnostic monitoring facilities built into the modems. Systems from other vendors, such as Avant-Garde Computing and Emcom, use independent monitors installed at the interface between the modem and the terminal equipment, between the modem and the network, or both. Devices that monitor both sides of the modem are said to "wrap around" the modem. Some vendors provide both alternatives: integral diagnostics in their own modems and wraparound devices for use with other vendors' modems.

The network management processor usually resides at a central site along with the network's host computer. Remote monitoring devices communicate with the network management system by one of two techniques. In the *mainstream* technique, favored by IBM and other vendors of host-based network management packages, the central-site unit (either the host processor or an independent network management processor) polls the remote devices (modems or diagnostic units) in a dedicated time slot over the main data channel. In the *sidestream* technique, favored by most modem vendors, the remote devices transmit diagnostic information asynchronously over a special, low-speed data channel frequency-division multiplexer onto the same facility as the main data channel. Some network management systems, usually marketed by test instrumentation vendors or network management system integrators, can support either mainstream or sidestream monitoring.

The mainstream technique lends itself well to end-to-end monitoring at the application level, since it is often directly in touch with the network's equipment, from host processor to modems and terminals. Host-based mainstream systems have a singular disadvantage: when the host goes down, the network control system goes down with it. One point of view holds that the network itself is useless without its host, but in many modern distributed processing systems the network can support a high level of activity even in the absence of a controlling host.

Modem/Multiplexer-Based Network Management

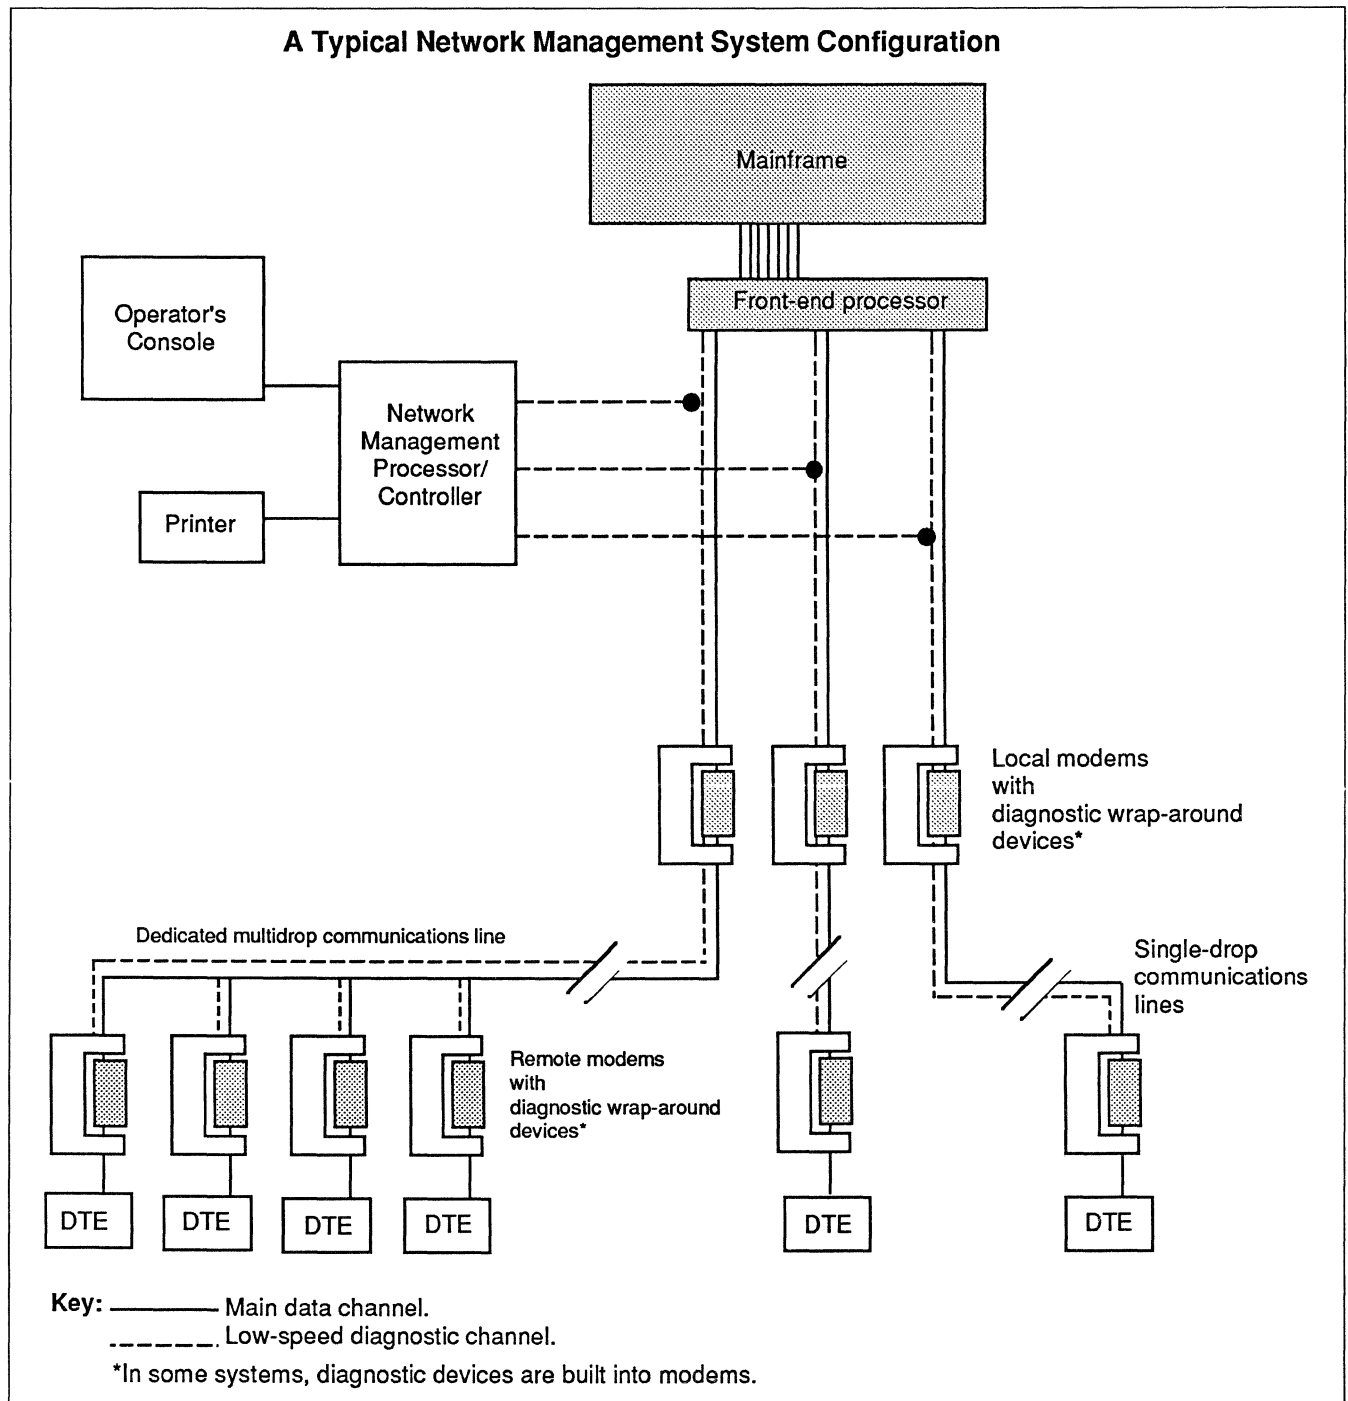


Figure 1. This diagram shows the overall configuration of a typical network management system. Network management components are shown as white boxes; main data network components are shaded. The system shown uses wraparound monitoring units at the local and remote modems, and employs a sidestream technique for reporting monitor information.

The sidestream systems are somewhat quicker to report modem, terminal, and communications line failures, since the monitors need not be polled in order to report trouble. The sidestream technique also offers a greater chance of survival over degrading com-

munications lines, since a low-speed signal is more likely to reach its destination intact over a noisy or hit-prone channel.

Modem/Multiplexer-Based Network Management

NETWORK MANAGEMENT FUNCTIONS

The basic functions one should expect from a modem- or mux-based network management system fall into three general categories: failure management, performance management, and configuration management. Early modem/mux-based network management systems generally offered only failure management, but more recent systems encompass all three. Any of these functions can serve both the operations level and the management and planning level.

Failure management breaks down into two functions: problem determination and system restoral. The problem determination facilities of a network management system can alert the operator to a failure, and on more sophisticated systems to a degrading condition, by using a scheme of alarms. When a local or remote monitor detects a problem, it sets an alarm at the operator station. Front-end processors, modems, switches, and multiplexers ordinarily provide a very primitive form of failure alarm, the disappearance of a "carrier detect" light on a supposedly active line. This is a negative alarm, and rather easy to ignore.

Network management systems typically provide a positive alarm signal, usually a message at the operator's console noting the type of device malfunctioning, the location, and the nature of the failure. Some systems set alarms for degrading conditions as well as absolute failures. Systems with color displays use a special color, usually red, for alarms. Color systems with multilevel alarms use green or blue for normal conditions, yellow for degrading conditions, and red for failures.

System restoral is a two-stage process. The user must first restore the network, then repair or replace the failed component. Some network management systems incorporate a fallback switching mechanism by which an operator can immediately bypass a malfunctioning device with a "hot," or ready-to-go-on-line, backup unit. For modems, multiplexers, terminals, and front-end processors, the backup units are available from a pool of spares, with one spare ready for a given number of active devices. For communications lines, the ready backup is usually the switched telephone network.

Most network management systems provide some automatic fallback switching, usually as an option. Some offer only manual fallback through a patch panel. Automatic switching can be either electromechanical or programmed. A number of switching configurations are possible. The A/B technique switches a single line between two devices, the malfunctioning unit and its one-for-one spare. One-by-N switching

switches a specified group of lines between a multi-line controller and its redundant backup, such as a spare front-end processor. In N-by-N and N-by-M techniques, any line within a specified group of lines can be switched among any of a given set of devices; the number of lines may be greater or less than the number of devices.

Fallback switching is a short-term answer to network problems. For long-term repair or replacement of faulty parts, network management systems provide two kinds of help. The first, available in most systems, is a facility for detailed testing to further isolate problems that have generated alarms. Some systems provide an integral test facility, while others merely provide access ports for the attachment of independent monitors and testing devices.

The second is a management device, the trouble ticket database. Paper trouble tickets have long been a mainstay of network and computer operations. Basically, a trouble ticket is an expanded version of a system log entry. It contains information on the date and time of a problem; the nature of the problem; the specific devices and facilities involved; any short-term actions taken to relieve the problem; the name of the operator who took the action; and space for recording follow-up information, such as visits from the vendor's maintenance staff, dates on which parts were returned for repair, serial numbers of spares installed, and the date of the problem's final resolution.

A trouble ticket database extends the logging and historical function of the trouble ticket to the management and planning level. The manager can call up reports on all trouble tickets that are outstanding, involve a certain subset of the network, are/are not recorded or resolved within a given period, and involve a specific device or a specific vendor. With such reports, a manager has an objective handle on such factors as the reliability of a given component, the promptness of a given vendor's field service, the reliability of a given operator, and the general proneness to failure of certain network segments.

Performance management deals with the up-and-running network from two viewpoints: response time and availability. Most network management systems measure response time at the local end, from the time the monitoring unit receives an "enter" or "end of transmission" signal from a given unit to the time it receives and passes a response back to that unit. Some systems can measure end-to-end response time at the remote unit. In either case, the network management system displays and records response time information and generates user-specified response time statistics for a particular end-user device; a par-

Modem/Multiplexer-Based Network Management

particular line; a particular subset of the network; or for the network as a whole, in realtime, or over a specified period. More sophisticated performance measurement systems can split the overall response time for any of those subdivisions into network time and computer response time.

Systems with color graphics consoles can display elaborate, multicolored network schematics with specific colors assigned to specific levels of response time, with blue or green for normal operation, yellow for degrading response time, and red for critically high response time. As with failure alarms, operators can use this information to reroute traffic and patch in additional resources.

Availability is a measure of actual network uptime, either as a whole or by segments. Availability statistics can include such measures as total hours available over time, average hours available within a period of time, and mean-time-between-failure.

With long-term response time and availability statistics, processed and formatted by the network management system, managers can establish current trends in network usage, predict future trends, and plan the assignment of resources for specific present and future locations and applications. Managers can also isolate and analyze chronic bottlenecks to specific areas and components. When an application overruns its allotted time, should management install more terminals and assign additional personnel or simply install faster communications? Response and availability information from a network management system provide objective tools to resolve the problem.

Configuration management involves both failure management and performance management, along with long-term planning of the network's topology and inventory. Some network management systems, in conjunction with the trouble ticket function, provide cost and depreciation information on the network's components. Most systems provide an inventory database with information on both active and spare parts. The network management system, used properly, gives managers objective information for decisions on purchasing and expansion.

THE MARKETPLACE

There are a variety of vendors supplying network management systems, including large modem and multiplexer manufacturers, communications test equipment specialists, and network management systems integrators. Each type of vendor supplies a different kind of network management system. The modem vendors consider network management to be

a value-added function atop their entire product lines. The test equipment vendors see network management systems as top-of-the-line integrated test systems. For the system integrators, network management systems are often the sole product. Each type of network management system has inherent advantages and restrictions.

The modem and multiplexer vendors, such as Memotec, AT&T/Paradyne, Codex, Racal-Milgo, and General Datacomm are strongest in end-to-end monitoring and alarming, and weaker in fallback provisions and long-term performance measurement. These vendors concentrate on selling modems and multiplexers and see the network management function chiefly as a way to sell more of their devices. In most cases, these systems will work only with remote devices from the same vendor and can "lock" a user into a single-vendor network. While these companies originally perceived their network management systems as a value-added function to top off their complete line of products, this attitude has changed. Realizing that a single vendor is unlikely to provide the answer for all network problems, vendors are offering systems that work in hybrid networks and with other vendors' products.

The increasing availability of end-to-end digital communications may soon affect the market for modem-based systems. Users whose plans include large-scale conversion to digital transmission may be reluctant to invest in a potentially obsolescent modem plant. Vendors are beginning to answer this objection by promising equivalent diagnostic support in future digital data service units (DSUs).

Integrated network management is another issue facing equipment vendors. Going from a single vendor approach to a multivendor network creates numerous problems. As a step towards integrated network management, vendors are promising support for third-party components through interfaces. For example, Racal-Milgo is changing its long term strategy in order to provide integrated network management systems. These systems and products are monitored and controlled from one site and will support other vendors' products.

Concerning integration, Jim Herman, of Northeast Consulting Resources, Inc. of Boston, sees two dimensions: product line and multivendor integration. Vendors who want to compete in the network management systems market need to have a game plan. First they must integrate existing products, provide a common user interface and access, and create an architecture for new products. Then, to compete in multivendor integration, vendors need to promote

Modem/Multiplexer-Based Network Management

their architectures as the base for integration, offer open interfaces, and provide common displays and databases.

Even though users do not yet have comprehensive, integrated network management tools, they are going forward with plans to upgrade their existing commu-

nications facilities or to build new networks. Many users feel multivendor networks are more viable than "one-vendor" solutions. Most industry analysts feel that it will take at least two more years before comprehensive, integrated network management systems will come to fruition. □

Network Management Systems for Data Communications

This report will help you to:

- Apply the concepts of monitoring, control, trouble management, resource tracking, and network planning to modem-controlled networks.
 - Examine the factors driving the development of network management systems.
 - Know what to expect from the next generation of network management systems.
-

In this report we will briefly summarize the five major aspects of network management: namely, monitoring, control, trouble management, resource tracking, and network planning. In a later section, a list of factors will be discussed that are currently driving the development of network management systems today.

Network management is discussed from the viewpoint of a telecommunications customer, rather than from the viewpoint of a vendor. Today's mixed vendor environment has created a number of challenges for both customer and vendor. Issues related to regulatory and organizational constraints are not covered, nor are management services that a vendor might be expected to provide to a customer.

FIVE ASPECTS OF NETWORK MANAGEMENT

Network Monitoring

Network monitoring involves the gathering of performance and diagnostic data, including alarms from the

This Datapro report is based on "Network Management Systems for Data Communications" by Andres C. Salazar, Philip J. Scarfo, and Robert J. Horn, Inffinet, Inc. © 1987 IEEE. Reprinted, with permission, from IEEE COMMUNICATIONS MAGAZINE, Vol. 25, No. 8, pp. 21-27, August 1987.

network, while the individual network components are in service. This data is expected to yield insight into developing problems that can be dealt with before they seriously affect network operation. Parameters in analog modem networks such as receive signal level and phase jitter level could indicate degrading line conditions if measurements are taken and analyzed periodically. Raw data from the network monitoring activity can also be used for trending analysis in the long-range network planning process.

Since network components can number as high as hundreds (if not thousands) of units, the gathering of the network data must be planned carefully. The protocol used to communicate with the diverse segments of a network must have the addressing capability and efficiency to handle heavy data traffic commensurate with network size. The volume of expected network status information must match real-time data acquisition capability of the network management machine. Hence, raw data should be carefully screened

Index to This Report	Page
Trends in Network Management Systems	203
Next Generation System Architecture 205
Expert Systems in Network Control 208

Network Management Systems for Data Communications

before being logged into a database. Mechanized entry into a database can cause performance problems if the transaction involves many index computations and file accesses.

A major problem in the monitoring of diverse network elements occurs when both hardware interfaces are different and disparity exists in the diagnostic data format. A network management user needs to access all network data from once console rather than have several control consoles for various network segments.

Automation of specific parametric measurements collected from network elements can be done on a user-programmed basis. Measurements taken while the network is on-line can yield one level of performance knowledge. Time of measurements and measure interval between replications are part of the automatic monitoring activity. A more detailed level of network performance measurements can be realized if links can be taken off-line and communication channels analyzed more thoroughly. These disruptive tests can be run during off-hours so as not to disrupt network operation.

Network Control

Upon determination that a network problem has occurred through the receipt of an alarm condition, the private network user will then wish to enter the network control mode of his/her equipment. In this mode, he/she hopes to further isolate or identify the individual network elements which may be causing the problem. This is done by performing disruptive tests such as end-to-end error runs or disabling a terminal due to its streaming behavior. Restoral techniques such as using dial-backup lines in order to reinstate failed communication links or activating hot-spare equipment in order to bypass failed units are an important part of network control (see Figure 1).

As discussed previously, the monitoring aspect of network management is highly automated and, in general, constitutes the mechanized accumulation of data from the network. The control aspect of network management relates to the human interactive response to the receipt of monitored data. Typical responses include running disruptive tests on a communications link in order to further characterize a reported fault and activation of service restoral equipment to work around the fault. These responses cause commands and in some cases formatted data to flow from the control site into the network. It is essential that the user be able to react quickly and

precisely to reported network faults. This real-time response capability is the essence of network control.

Trouble or Problem Management

Large networks often have a significant number of reported problems which need to be recorded and tracked before final resolution. The keyboard entry of this information occurs at a trouble desk where dispatching of personnel to resolve the problem is often done. Summary or status reports need to be generated periodically to ensure that the number of problems has not gotten out of hand. Data concerning reported network problems can be used to generate time-to-repair or out-of-service averages. These parameters are often closely monitored for trending information. The trouble desk also needs to have the ability to associate different levels of criticality to problems not only at the time of first report but later when non-resolution may require escalation of attention.

Since keyboard data entry and report generation constitute the major functions of trouble management, the user interface for executing these functions becomes important. System response time for entering data and creating reports is a critical parameter. The database used for this function of network management must have the tools necessary to generate customized reports which relate one set of data fields in one application differently from a set of fields in another report application. Figure 2 illustrates the different activities associated with trouble management.

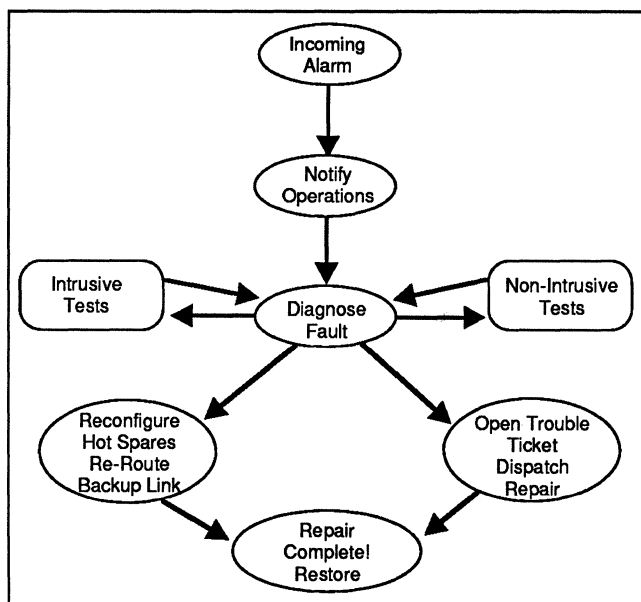


Figure 1. Alarm handling.

Network Management Systems for Data Communications

Resource Tracking

A private network consists of a large number of elements which differ widely in type, vendor, location, service level and functional criticality. There is a recognized need to inventory all the network elements for the purpose of asset control in addition to expediting maintenance and service restoral procedures. This need for asset control has increased dramatically in recent years as private network managers have been driven to maintaining and servicing their own networks. Self reliance in network management often requires knowing where equipment spares are located and who is able to complete the repair procedure. The diversity of network equipment which often parallels an equal diversity of suppliers or service bureaus has led many network managers to require resource tracking capability from their network command center.

Since the resource tracking function of network management is again data base intensive, keyboard data entry and report generation again become important user interface issues. Data base updates should be done quickly with minimum system transaction delay. As in trouble management, customized reports are an important feature to users whose needs for data relationships change from time to time.

Network Planning

Every network of data communications equipment has a set of performance measurement parameters whose values provide the network manager a network service index. These parameters include measure-

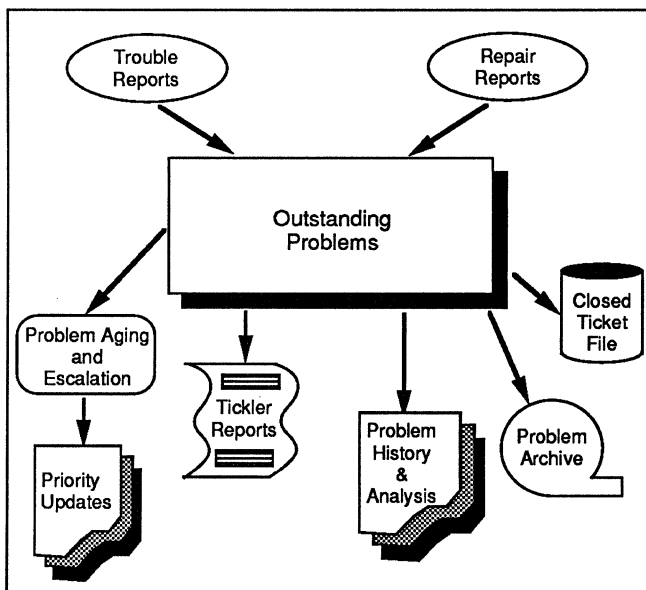


Figure 2. Problem management.




Figure 3. Typical '60s telecommunications network.

ments such as number of terminals out of service for more than an hour, mean time to repair for last ten service calls, etc. In the end, downtime costs money as does restoral equipment. Hence, tradeoffs need to be made on the basis of (a) historical trouble data (b) inventory and service costs (c) business changes (d) line tariff changes. Network design can proceed on an optimized basis only if reliable data can be kept on the network performance parameters. Network managers find that flexible summary reports on resources and trouble tickets are invaluable in the network design process.

TRENDS IN NETWORK MANAGEMENT SYSTEMS

Integration—The New Network Management Problem

In today's networks the underlying communications equipment has added another large layer of complexity. Networks previously were internally uniform. They were mostly analog networks, predominantly serving one major application, with one control center, and with one major network services provider: AT&T. (See Figure 3). Today's larger networks (see Figure 4) will involve:

—multiple network service vendors for the same technology

—multiple transmission technologies:

- analog facilities
- digital facilities
- T1 and microwave facilities
- satellite links
- Local Area Networks

—additional transmission processing functionality:

- multiplexors
- packet-switching
- remote polling

Network Management Systems for Data Communications

- multiplexors
- packet-switching
- remote polling

—mixed purpose systems—voice and data combined

- ISDN

Each of these facilities may have control capabilities built in by their manufacturer. This leads to a wild proliferation of consoles at the network control center. Network managers now have to juggle several terminals and use the capabilities of several vendors' equipment to isolate and solve a single problem. A central control system that evaluates and diagnoses the entire network often does not exist. In its place are many separate and uncoordinated control systems for various network segments.

The control center is also being asked to do more than find and fix line problems. Now its job includes monitoring remote terminal equipment and maintaining system performance standards. These jobs add still more to the list of tasks and measurements to be taken and evaluated. The load for this is substantial. For example, just recording the performance statistics for 500 terminals on an hourly basis can consume between 500,000 and 1,000,000 bytes per day. All this data requires processing to collect, to transmit, to evaluate, and to store and retrieve.

Transmission of data is no longer "free" either. With the shared facilities like packet switching every byte sent cost money. When the network was built of

leased lines and the costs were traffic independent it was reasonable to send everything back to the center. Now the cost may be substantial when some of the links are priced on a traffic basis.

Very large networks also are evolving with multiple control centers. Some corporations find it preferable to have regional control centers which have responsibility for some portion of the network. This may be limited to only troubleshooting related control access, or it may be a complete backup control center that is ready to step in and take over control in the event of a major failure at the primary center. Whichever organization is used, this means that the communications equipment in the network must now be designed to have several masters. It must keep track of which site(s) should be sent data and alarms, and still be responsible to any authorized controller for testing.

Another major problem today deals with the potential need to integrate local area networks (LANs) into the wide area Network Management System. Many of today's LANs have been expanding steadily since their introduction a few years ago. As they expand, network management becomes an important consideration to controlling them. A network with several hundred nodes offers many of the same problems to the network manager of a wide area network. Alarms, network performance, and configuration control are but a few of the common problems presented.

Unfortunately, most of our LANs today have not been designed with network management in mind. Consequently, little or no diagnostic information is available to signal points of failure or potential bottlenecks in the network.

Configuration changes must be better controlled and regulated as networks grow. The addition or deletion of nodes on the network must be known to the network controller to maintain proper system availability.

In short, we have a number of new network management problems confronting today's network management designers. The "integration" battle represents one of the major new challenges ahead.

Second Generation Systems

Many of today's second generation Network Management Systems are products of the minicomputer architectures of the last decade. Predecessors to the new 32-bit super-micro computer, these 16-bit computers delivered compute speeds of 0.5 to 1 MIPS. Yet despite their rated performance, many of these first and

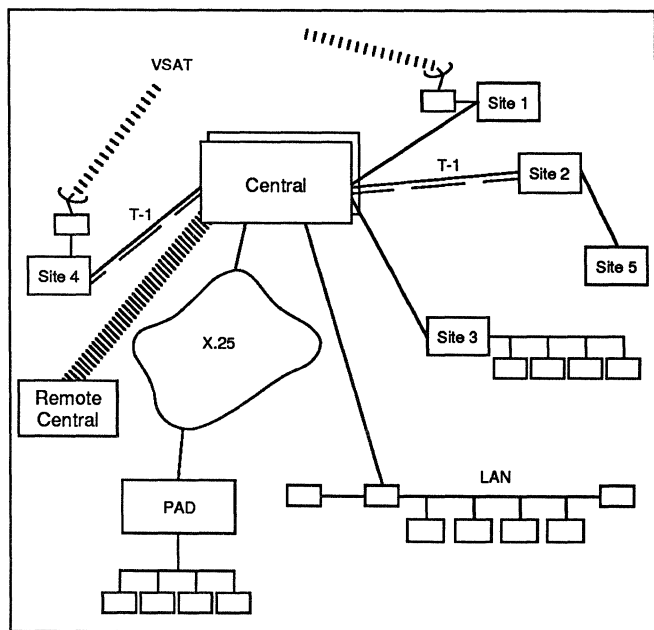


Figure 4. Typical '80s telecommunications network.

Network Management Systems for Data Communications

second generation systems are a poor match to the real-time demands of today's larger telecommunications network. Even super-minis like the VAX 11/780, rated at about 1 MIPS, are underpowered for today's large hybrid network.

User complaints about system performance and the inherent lack of flexibility provided by a single processor system have limited the success of many of these systems. A single processor system (Figure 5) also offers limited upgrade potential for the growing telecommunications network. Oftentimes the only upgrade path is one requiring a wholesale exchange of the system. Both capital and data investment are lost.

Beyond the flexibility issue, the absence of industry standards for network control diagnostics has had a major effect on these early system designs. Proprietary system architectures have been adopted by most data communication vendors. Understandably, their primary focus has been on the development of data communication products and not system based products. The result is decreased flexibility for both network planners and operators.

With proprietary hardware or designs which were based on a computer vendor's closed system concept, it has become increasingly difficult for data communication vendors to take advantage of the tremendous strides in computer systems technology. Many systems have applications software which is not easily transportable to these cheaper and faster devices.

New processors like the Motorola's 68020 and Intel's 80386 already offer computational speeds in excess of 2 MIPS. This is better than twice the speed of these early minicomputer systems and supermini systems like the VAX 11/780. Yet they are available from multiple vendors today at only a fraction of the cost

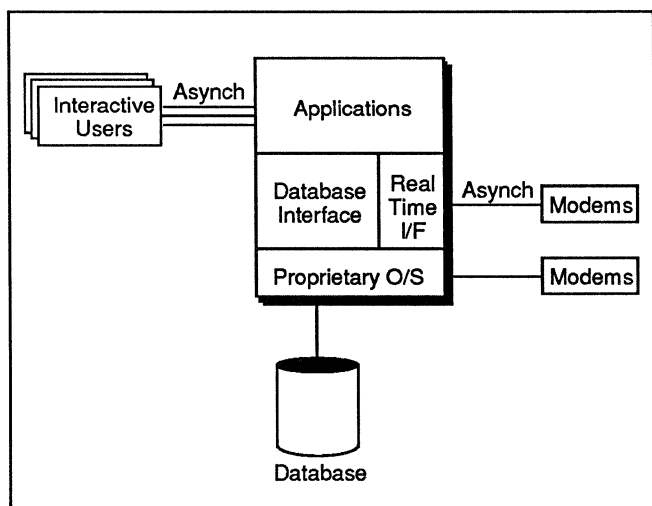


Figure 5. Single processor system.

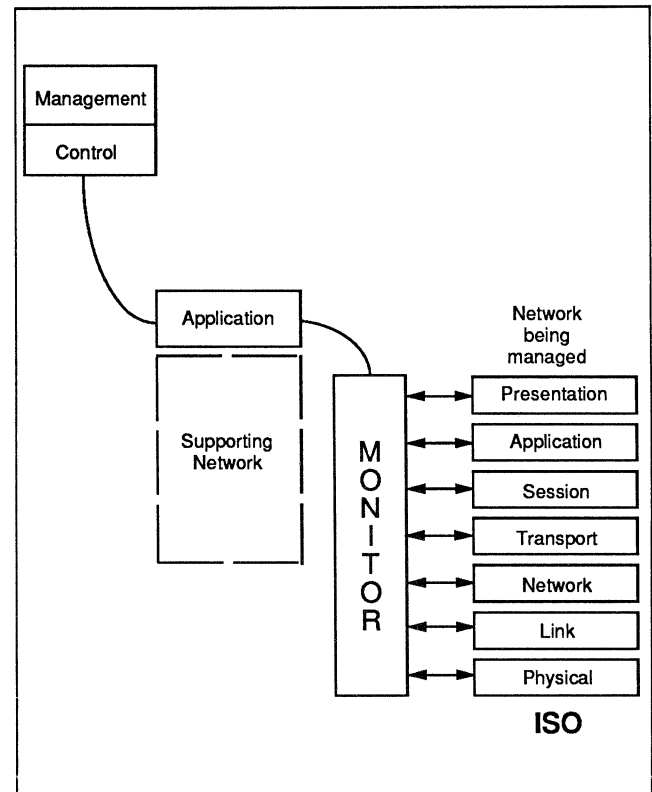


Figure 6. Network management architecture.

their super-mini counterparts. Lower cost memory, mass storage, and high speed local area networks have also changed the design strategies of network system designers, as more distributed architectures become a cost-effective alternative.

Next Generation System Architecture

Network management is not structured in conformance with the ISO reference model. It is an application and is structured to meet its internal needs. However, since it is managing a data communications network which probably does match the ISO structure, it will be very strongly influenced by the ISO reference model.

The network management architecture is shown in Figure 6. There are three major components:

- The management layer, which focuses on the external and internal organizational issues of
- Trouble or problem management
- Resource tracking
- Network planning

Network Management Systems for Data Communications

These tend to be highly interrelated and usually share a common database.

—The network control layer

—The network monitoring layer

Network monitoring has interfaces to each of the layers of the data communications network. In a fully configured monitoring system there would be a monitoring interface at every termination of each layer. This means there must be a measurement point at each access point into a physical media, each link access point, each network termination, etc. Costs may restrict the monitoring to just the critical points.

The monitoring layer is connected to the control layer via the application layer services of some support network. This is usually necessary because monitoring points occur throughout the network. One serious problem that results from this configuration is that the support network may also be part of the network under control. In this situation application level commands may temporarily disrupt their own underlying network. This imposes the extra constraint on network design that an application layer connection survive complete (temporary) loss of the lower layers. Some networks cannot do this, and for these the support network must be separate from the network being controlled.

The nature of different telecommunications technologies controls the internal structure of the monitoring layer. The kinds of monitoring and controlling activities that are appropriate to a LAN differ from those for a WAN. Different technologies lead to domains where each individual domain corresponds to a particular set of layers and interfaces that are tightly coupled. For example, 802 LANs, X.25 networks, and T1 networks would each be a unique domain even though they provide services that appear in the same ISO layer. The monitoring and control interfaces to these domains are the focus of much of the current standardization efforts in network management.

The network layer becomes a collection of individual domain monitoring modules, each focused on only a few layers and supporting a specific technology. These modules may co-exist in the same processor, but they operate autonomously.

The network control layer comprises two sublayers. The lower sublayer is a collection of domain control modules in a one to one relationship with the monitoring domains. Each of these handles the control actions that are appropriate for their associated technology. The upper sublayer handles control actions that require coordination of multiple domains. This

layer is only rarely implemented in current systems. It is usually handled directly by the network control personnel.

The network management layer comprises three different views of an integrated network management database. These correspond to each of the major management functions. The control layer provides continuous updates to the database to maintain a current view of the network. These updates are used by the network managers to determine both reconfiguration and organizational actions. Network managers also update the database to track and control their organizational activities.

In Figure 6, both the network management and control are shown located on a single host. This has the advantages of rapid interaction between management and control, and the advantage of sharing a single network database. It has the disadvantages of imposing maximal traffic loads on the support network, introducing network delays for control to monitoring layer interactions, and requiring a large central facility. A further disadvantage is that it now becomes a single site failure risk for the entire network. These problems can be reversed by shifting the control modules to the other side of the support network and placing both the control and monitoring activities for a domain into a domain processor. This has advantages when there are serious network problems, because local control and monitoring facilities will survive even serious network failures.

The management layer may also reside on either a single host or a distributed system. The tight coupling of the user tasks and the database make distributed operation more difficult. Reliable performance is most critical when the network is failing, and this is the environment that poses the most problems for current distributed databases.

Next Generation Designs

A third generation network management system design can exploit the economics available in today's technology and utilize a multi-processor approach versus a single processor architecture. This provides both greater flexibility and system performance over time.

Conformance to emerging computer industry standards is also a must today. As computer vendors move towards a more "open systems" philosophy to protect their own investments, it becomes important for third party developers to take advantage of the flexibility and performance enhancements which emerge from these architectures.

Network Management Systems for Data Communications

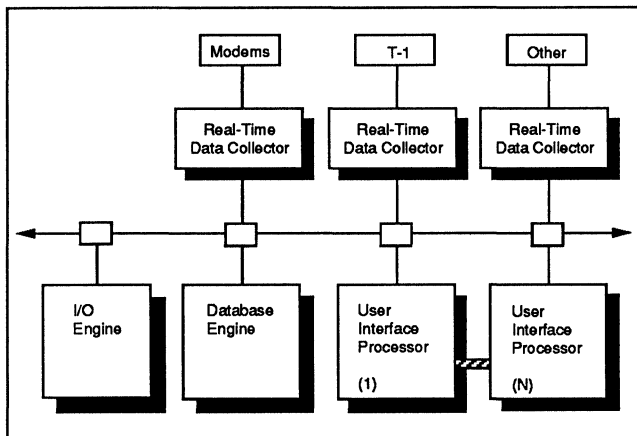


Figure 7. Multi-processor system.

New chip technologies, parallel processors, and new RISC architectures will provide even more impressive price-performance ratios in the future. Applications software which is designed to take advantage of these changes will emerge ahead of the pack.

Systems designed around de facto industry standards like UNIX, Ethernet, TCP/IP, and NFS may offer greater capability and provide the greatest protection against product obsolescence. In fact, in the data communications arena, a de facto standard may be more important than an official standard committee approval, since installed base may be the significant design consideration.

Distributed Network Architecture

A distributed network architecture which utilizes several processors working in concert can deliver superior performance and greater flexibility. Horsepower can be added where needed rather than require the replacement of an entire central processor. The ability to couple intelligent user or network processors with a network management processor on a high speed LAN backbone provides an incremental approach to upgradeability (Figure 7).

One example of a distributed architecture design might have a processor which serves as a "database" engine or central file server on an Ethernet backbone with several client "user" processors and "network interface" processors. Common information or files could be shared across this network yet the network management database itself would remain centralized under the control of the database engine. The user's investment in network data, problem management history, or other historical information can be preserved.

Each of the "engines" could be sized to the application requirements or added to the network as system requirements change. A separate dedicated network interface processor would be assigned to each class of data collection device on a hybrid network. This allows the network planner to plan network service expansion by simply adding a new front end data collector when needed.

Complex networks supporting multiple network domains, i.e., analog and digital devices, private X.25 or T-1 products, could then be more easily integrated into a single network management system. Software modifications would be limited to those front end data collection requirements of the new device alone.

Finally, as application and user demands increase, intelligent workstation processors can be added to the network control capabilities. Applications requiring high speed graphics processing or rapid command and control response can be better served by an intelligent workstation processor able to operate independently from the central data manager. These user

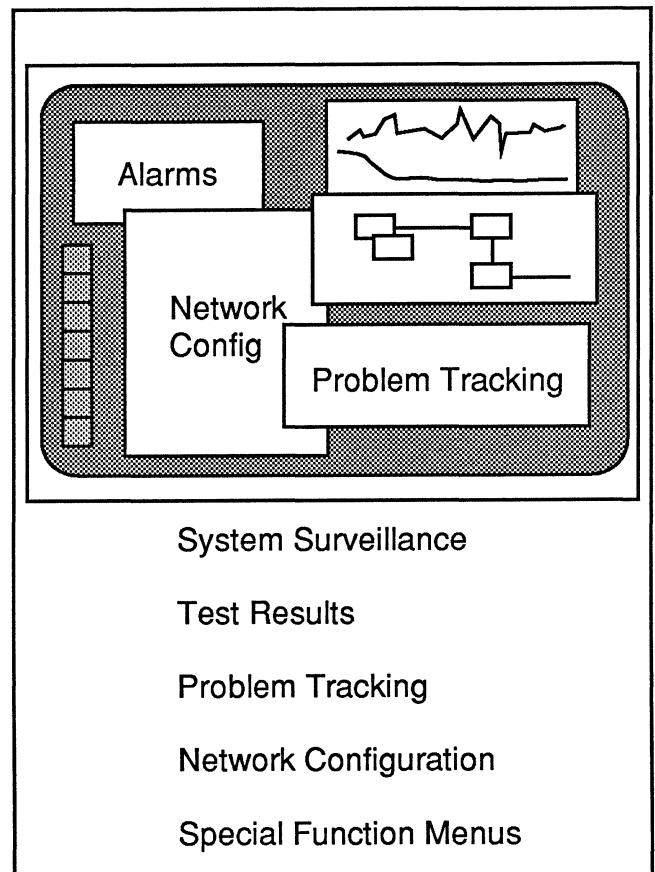


Figure 8. Multi-window environment.

Network Management Systems for Data Communications

User Interface Trends

The use of workstation technologies is changing the user interface from a text and forms orientation to a multi-window, graphics environment. The problem analysis activity calls for examination of results from multiple test and often from multiple domains. The windowing environments permit the simultaneous presentation of this information. There remain some areas where text is the appropriate format. Configuration details, field service addresses and phone numbers, and exact measurement results need the precision of text. But with the addition of graphics for displaying interconnections, trends of measurements, and status monitoring, the added power of human pattern recognition can be exploited. The new workstation technologies provide the processing and display power needed to present both the graphics and text in a timely and usable fashion.

The resulting systems present a much more dynamic image to the user. Routine surveillance is very graphically oriented, using color keys to highlight problems. Problem diagnosis and problem tracking activities take place in windows so that key surveillance monitoring can continue in parallel (Figure 8). The network controllers have a wealth of information available on command, presented in the format most suitable for problem management. This information is called up and dismissed from the screen as the operator commands. These are all geared to reduce the problem diagnosis and repair time.

Early examples of this can be found in the new products being delivered by network control vendors. These have not yet made a significant penetration into the market and they will be modified based on the feedback received from the customers of these early shipments. The concept of using graphics is widely accepted. There is no consensus yet on what presentations are most effective.

Expert Systems in Network Control

The computers themselves are also taking a more active role in network management. The greatest progress has been in:

1. Problem diagnosis—Expert systems are in use to speed the recognition and analysis of network faults. Instead of simply recognizing the fault and relying on the operator for diagnosis, systems are analyzing the fault indications and network status. They can then present not only the fault, but also a composite fault diagnosis and then suggest corrective action.
2. Network configuration—The great many elements and rules associated with configuring communications lines, modems, multiplexors, etc. presents a major burden to the network designer. Expert systems are also being employed to aid in the configuring of complex networks. There will be increasing use of expert systems and artificial intelligence techniques to handle the routine surveillance and diagnosis activities and to assist in other network management activities. □

The information provided in this tab is researched and written by NBI/Datapro. NBI provides the very latest intelligence on the telecommunications market and corporate strategies. NBI's research and industry analysis includes detailed information on key telecommunications markets—both domestic and international.

The reports in this tab will help you understand the strategies of the leading network management vendors. These reports identify critical strengths and weaknesses of carriers and equipment suppliers currently in the marketplace.

Achievement of your goals requires many short-term considerations, but long-term relationships with vendors are key: you must understand what direction your carrier is taking its network, and where your equipment vendor is taking its architecture three, five, and even ten years from now. The reports in this tab are designed to do just that.

NBI/Datapro is the leading supplier of telecommunications market research in the world. For further information about our products and full-support services, please contact us at:

Northern Business Information/Datapro
157 Chambers Street
New York, NY 10007
Telephone: (212) 732-0775
Fax: (212) 233-6233

Strategies of Major Network Management Vendors

This report will help you to:

- Anticipate trends in the network management market into the 1990s.
- Examine the network management strategies of key vendors, including IBM, AT&T, Digital Equipment, and Northern Telecom.
- Evaluate which strategies will be successful in the coming decade.

For the customer, network management means multivendor solutions that reduce costs. For telecom equipment and service providers, network management means a radical shift in strategy.

MARKET FORECAST AND ANALYSIS

Telecom and computer equipment vendors bank on network management products and services to inject some new life into hardware sales over the next five years. Success for all will require quick acceptance and implementation of standards—traditional proprietary protocols that lock out competition will not work in this market.

Figure 1 shows the network management equipment market for 1987 through 1993. Projections for public network management equipment for carriers (or operations support systems—OSSs) and for private network equipment are shown. Services are not included.

We make the following assumptions about the network management market:

- OSI standards for private network and OSS applications are sufficiently developed and accepted by the end of 1990 to permit unimpeded product development;

- OSI-based network management systems (NMSs) achieve equal market share with SNA-based systems by the end of the forecast period;
- AT&T continues to dominate the OSS market, in contrast to fierce competition in the private network NMS arena.

This report reviews the strategies of the key participants in network management.

IBM

For IBM, the big question is how far NetView and NetView/PC will extend from SNA private line networks into AT&T's domain, the public switched network. When SNA can be fully supported on ISDN (i.e., 23B+D, not 2B+D), IBM will be positioned to tap much of the public switched network's potential to add value.

IBM's future lies in value-added service, and by 1992 it will begin its big push into communications and transform itself into a services-oriented company. IBM is now laying the groundwork. IBM's PS/2 has nearly the same processing power as some early versions of the System/370, and by 1992 there is good reason to expect that many users will have /370 power on their desks. By this time, distributed database management and effective peer-to-peer communications will be crucial—two capabilities difficult or

Strategies of Major Network Management Vendors




Figure 1. Network management equipment market, 1987 through 1993.

impossible to implement on private line networks. This is the market IBM is going after in telecom. By its sale of Rolm, IBM confirms hardware is no longer its primary interest.

Thus unshackled, IBM is now free to team up with many telecom hardware vendors to make voice network management a reality for NetView. Such multi-vendor arrangements are essential if IBM is to keep management of voice networks from being dominated by AT&T and Northern Telecom and to check the ascendancy of NetView substitutes, such as Cincom System's Net/Master.

AT&T

AT&T is not standing still. The company has announced the broadest network management platform imaginable, its Unified Network Management Architecture (UNMA). The company is positioning the architecture for both public and private networks. AT&T has already published its proposal for a standard network management protocol (called "NMP") for computers, switches, and transmission equipment. Not surprisingly, AT&T is staging its attack from the public network. For example, NetPartner, one of the early UNMA products introduced in September 1988, provides a glimpse of AT&T's future strategy.

The NetPartner system uses 3B2/600 series computers, the Datakit circuit/packet switch, and Sun-3 workstations. NetPartner's first applications will be for testing, surveillance, traffic reports, reconfigurations, and billing. As usual, AT&T markets NetPartner as an "ISDN" product, but it does not require

ISDN lines. NetPartner can support 5 to 10 Centrex customers and up to 40,000 analog, digital, or 2B+D lines.

An agreement between AT&T and Cincom Systems, Inc. might pose an even bigger challenge to IBM. (For more information on this agreement and on AT&T's ACCUMASTER Integrator product, see "AT&T Unified Network Management Architecture," Report NM40-313-101.) Cincom markets Net/Master, the only real current threat to NetView. If IBM's customers view the sale of Rolm as an abandonment of voice networks, AT&T and Cincom might together fill the breach.

Digital Equipment Corporation

In September 1988, Digital released a comprehensive network management plan called Enterprise Management Architecture (EMA). EMA, a highly flexible architecture, supports multivendor connectivity via OSI's Common Management Information Protocol (CMIP) or through proprietary interfaces. Digital's EMA is well adapted for distributed processing environments. Seven vendors, including specialists in modem, bridge, and T1 technology, have teamed with Digital in the EMA development effort. Digital has not yet announced specific EMA products; while progress appears slow, we expect to see an EMA developer's kit by third-quarter 1989 and a first round of products in late 1990. For more information on EMA, see "Digital Equipment Corporation Enterprise Management Architecture," Report NM40-325-101.

Strategies of Major Network Management Vendors

Northern Telecom

Northern Telecom has a limited range of NMS products—and limited sales—given its large role in telecom but is moving on several fronts to extend its network management product line. In 1988, the company introduced “Meridian Network Control” for its PBX networks. A previously introduced product family, the Dynamic Network Controllers, targeted both MIS and BOCs for private network and local Centrex services. Designed for Bell Canada, the DNCs have never gotten off the ground in the U.S. The “Digital Facility Management System,” introduced in 1984, has been sold to the seven BOCs for digital carrier management.

A joint venture between Northern and Hewlett-Packard, formed in 1988, provides systems integration for large corporate networks. Hewlett-Packard has already developed network management for private networks, and in 1988, introduced a family of products called OpenView. OpenView Windows, a graphical interface, is currently the most notable member of the product family. (For more information on HP OpenView, see “OpenView’s Architectural Models,” Report NM40-452-101.) Both HP and Northern Telecom are looking for hardware pullthrough from their joint venture.

Others

Carriers other than AT&T are also entering the fray.

U S WEST Network Systems, Inc. (WSNI), a subsidiary of U S WEST, started on the private side with a

PC-based product for SNA networks. (For more information about this product, see “U S WEST Network Systems, Inc., NetCenter Products,” Report CMS60-950-101 in *Datapro Reports on Communications Software*.) U S WEST recently rolled out another network management system for Centrex and PBX. Now the company is trying to bring the two together. An interesting aside: U S WEST has an OEM agreement with TSB International, a Canadian company which makes a product that collects call detail reports for Rolm PBXs (for which IBM has already built a NetView/PC interface).

In October 1988, MCI bundled NetView with its virtual private network service, called Prism. In a similar fashion, Tymnet is offering network management capabilities bundled with its packet-switching services. Customers may be attracted by such a “holistic” solution to their problems.

THE FIVE-YEAR OUTLOOK

During the next five years, success will go to the suppliers that “sell into the cost reduction curve.” Network management reduces operating costs. As an added benefit, network management opens up the customer’s installed base to competing products, anathema to classic vendor marketing strategies.

Vendors have little choice but to let customers have their way. Large users and carriers have forced the biggest telecom and computer manufacturers to offer integrated management alongside integrated communications. Each step toward integrated management is a small victory for users and carriers alike. □

User Evaluations of Vendor Offerings

This report will help you to:

- Anticipate the dramatic increase in network management expenditures over the next three years.
- Evaluate network management purchase criteria among telecom managers.
- Identify network management buying intentions.

In March, NBI/Datapro, in conjunction with McGraw-Hill Research, completed a survey of telecommunications and MIS managers regarding their plans for network management and general network use. The survey intent was to determine current and planned use of network management systems and services (NMS) by a cross section of American organizations.

The survey quantifies network management buying intentions with detailed analyses by vertical market, company size, current use of network management systems, and current stage in the NMS purchase process. The information contained in this report has been obtained from sources we believe to be reliable, but neither its completeness nor accuracy can be guaranteed. Opinions expressed are based on our interpretation of available information.

NETWORK MANAGEMENT STANDINGS

Seventy-seven percent of the respondents ranked AT&T as "one of the best" or "above average" in NMS capability. Nearly 60 percent ranked IBM tops in the field, and between 40 percent and 50 percent ranked Digital Equipment Corporation, Northern Telecom, MCI, Hewlett-Packard, and the local exchange companies (LECs) first.

Surprisingly, current NetView users ranked AT&T at the top, although they did rank IBM higher than did nonusers. MCI, Hewlett-Packard, and the LECs fared considerably poorer with NetView than with other users. NET, however, fared much better with NetView users, probably because of NET's association with IBM.

Twenty-five percent of the respondents see in-house development as the primary source of network solutions over the next three years. A surprising number of companies want their interexchange carriers (IXCs) to provide or improve network management; the network control features now available from AT&T (Accumaster) and MCI (MCI View) should be well received.

NetView will account for 17 percent of the respondents "primary source" of network management systems; OSI-based solutions will account for 8 percent. Of the 41 respondents in the "actively reviewing" or "still defining needs" category, 5 percent said they would buy NetView.

Only 6 percent of respondents said they plan to use facilities management services during the next three years. However, 12 percent of the NetView users, 11 percent of government organizations, and 8 percent of the companies with revenues over \$100 million said they would contract for facilities management services.

User Evaluations of Vendor Offerings




Figure 1. 1991 network management purchasing plans (percentage of respondents who will increase buying volume by 1991.)

The most common types of NMS currently in place include intelligent network elements (e.g., PBXs, T1 multiplexers, modems), alarm systems, and test equipment. Non-NMS users plan to buy this type of equipment (see Figure 1). Over 50 percent of current NetView, UNMA, intelligent elements, test equipment, and OSI software users plan to increase their spending on NMS.

The survey indicated significant interest in systems integrators for network management and other purchases. Only 21 of all respondents place a high value on one-stop shopping; but only 12 percent of NetView users ranked it high. Marketing opportunities exist here; NetView users will buy pieces of networks.

The survey also showed that buyers want responsiveness, reliability, and good service from their network management suppliers; users in the buying process ranked these vendor qualities the highest. Low prices and technical innovation ranked much lower in the buying decision (see Figure 2).

Profile of NetView Users

NetView users, while committed to the NetView approach, by no means rule out incorporating other vendors into their networks. Indeed, NetView users are less likely to place a high value on single-sourcing than non-NetView users.

The survey revealed a number of distinctions between respondents who now have NetView and those who do not. NetView users:

- are heavy investors in all telecom elements, except low-speed modems;
- will not increase purchases of CDR, centrex, or low-speed modems;

- will buy more fiber optics, T1 multiplexers, voice-data PBXs, and voice mail than non-NetView users;
- plan to use CLASS, CO-LAN, digital centrex, information gateways, ISDN, and VPNs to a greater extent than non-NetView users;
- have higher communications operating budgets than non-NetView users now, but capital equipment budgets are the same;
- about half will find future network solution sources with NetView; the other half divides equally between in-house solutions and OSI-based NMS software;
- will increase usage and buying volume of NetView (64 percent); only 5 percent of those without NetView plan to purchase it;
- do not plan to use facilities management and systems integration, but do plan to use carrier-provided NMS, more so than non-NetView users;
- rate "single-sourcing ability" and "personalized customer service" less important than non-NetView users;
- rate NET better than non-NetView users and rate Hewlett-Packard, Unisys, and LECs worse than non-NetView users; both rate IBM and AT&T about the same.

BUDGET FORECAST

The survey showed an average increase of 7.7 percent in communications capital budgets by 1991. Companies with revenue over \$100 million showed an average increase of 7.9 percent in their communications capital budgets. NMS users' budgets will increase an average 9.3 percent; NetView users' budgets will rise




Figure 2. Network management purchase criteria (percentage of respondents who consider these factors extremely important when specifying vendors for network management systems and services).

User Evaluations of Vendor Offerings

10.2 percent; and NMS buyers will have 9.7 percent more to spend. Capital budget increases varied less than 3 percent among the groups. Interestingly, respondents placing a high value on single-sourcing reported plans for an average 10.2 percent rise in capital expenditures.

Network management customers plan to increase their expenditures over the next two years. By 1991, average expenditures for NMS will increase 3 percent (to 8 percent of communications capital budgets). In four categories of respondents, NMS expenditures will rise higher than the average 8 percent of budget. In companies with over \$100 million in sales, expenditures will rise from 7.9 percent to 11.1 percent of communications budgets. Companies with network management systems in place will increase budgets from 7.8 percent to 10.4 percent. NetView users' expenditures will rise from 6.3 percent to 8.9 percent. Thirty-three percent of the respondents are "currently buying" or "committed to implementing" NMS. Those respondents plan to increase their NMS spending from 8.5 percent to 12.7 percent of communications budgets.

The respondents' operating budgets will increase an average 8 percent by 1991. But the hot prospects for NMS vendors will increase budgets from 9 percent to 12.5 percent. Companies with over \$100 million in revenues will average a 9.1 percent increase in operating budgets; current NMS users will average a 9.7 percent increase; current NetView users will average a 12.5 percent increase; and companies in the NMS "buying mode" will average an 11.0 percent increase.

NETWORK USAGE FORECAST

The voice mail market shows the most activity: 24 percent of nonusers said they will buy voice mail over the next three years. Nonusers also expressed strong interest in communications and network software, high-speed modems, LANs, and T1 multiplexers.

Expect high-speed service demand to increase sharply during the next three years. The survey asked respondents how their telecommunications operating budgets were spent by service type in 1988, and what they expect the distribution to be in 1991. The greatest increases came in T1/DS1, T3/DS3, switched-56, and, to a lesser extent, facsimile communications. T3/DS3 demand is strongest among transportation and communications companies, the financial industry, and NetView users. T1/DS1 demand is concentrated among large companies and NetView users. Surprisingly, fax accounts for a larger share of small company budgets.




Figure 3. Vendor preferences (respondents evaluating network management capability of the companies shown as "one of the best" or "above average").

The survey also asked about interest in newer/advanced services. Virtual private networks (VPNs) garnered the most interest. Over 30 percent of the respondents said they plan to use VPN; 53 percent of large companies and a majority (72 percent) of NetView users plan to use VPN. A greater than average number of NetView users want IXCs to provide network management.

Two thirds of the respondents indicated that they plan to use dedicated fax networks. In light of the demand, other carriers are likely to mimic MCI's new fax network service.

Although many respondents stated an unfamiliarity with ISDN, 18 percent said they would use the basic rate interface, and 12 percent said they would use the primary rate interface. Most respondents indicated a need for better marketing before they would buy ISDN. The percentage of respondents planning to use CLASS and CO-LAN was low, but again many respondents said they were unfamiliar with those services.

Generally, the response to centrex was anemic, although, not surprisingly, best among current centrex users. Centrex users showed a strong interest in advanced centrex features. Telecommunications managers appear to be intrigued by two features which can dramatically improve network control: voice mail and ACD behind centrex.

CONCLUSIONS

Telecommunications and MIS managers will increase their network management expenditures between now and 1991.

Despite NetView's current strong position, IBM faces stiff competition in the network management market. Niche players and newcomers to the market have a good shot at cracking the IBM shops with their own NMS: buyers are not looking for single-source solu-

User Evaluations of Vendor Offerings

tions. In-house development and IXC services comprise the most notable alternatives sought by users.

For IXCs, opportunities abound. A strong interest in IXC solutions to network management problems exists even among NetView users.

Although many companies plan to implement single-point network control, they will continue to invest heavily in intelligent network elements, such as PBXs, T1 multiplexers, and other devices with built-in network management.

The largest category of users comprises those "still defining needs." Not surprisingly, service, responsiveness, vendor reliability, and after-sales technical support are the most important buying criteria. Generally, survey respondents placed a low value on single-sourcing, technical innovation, and price.

For more information, see NBI/Datapro Survey of Network Management Markets: 1989 Edition, May 1989. □

AT&T Network Management Strategy

This report will help you to:

- See how AT&T built its network management product strategy around its Unified Network Management Architecture (UNMA).
- Evaluate AT&T's network management strategy with regard to the company's current structure.
- Understand how the effects of divestiture will impact AT&T's network management product/service strategy.

AT&T is offering leading edge network management technology with which it plans to attack IBM's networking weaknesses and make big inroads into MIS markets. However, AT&T may not have the account leverage needed to persuade large companies to buy into its networking philosophy and methods.

Indeed, it appears that IBM and AT&T will have their greatest conflicts in the network management business. AT&T's ACCUMASTER products are designed to go head-to-head with IBM's NetView. Further, IBM and AT&T both want to enter the facilities management/systems integration business in a big way. For example, AT&T is teaming up with Computer Sciences Corp. (CSC) to compete with IBM,

MARKET POSITION

AT&T literally owns the network management business among the local exchange carriers (LECs), where network management is a mature activity—the Bell Operating Companies (BOCs) use AT&T operational support systems (OSSs) almost exclusively. AT&T is up against feisty competitors, such as Digital Equipment, that are designing a broad range of OSS products to cut into AT&T's business. Already well established at most carriers, Digital has enough control over customer buying decisions to seriously impact AT&T's business.

Among large corporations, however, where network management is a new process and all the rage just now, AT&T faces IBM, which has a lock on most large MIS accounts. In this end of the business, moreover, although AT&T's products are great—its sales operations are sluggish and ill suited to the sophisticated needs of corporate MIS.




Figure 1. AT&T network management sales, 1986 to 1993 (\$ in millions).

AT&T Network Management Strategy

MCI, and GM's Electronic Data Systems (EDS) to bid for the management of Merrill Lynch's telecom resources.

In its fight with IBM, AT&T is trying to leverage the power of its UNIX technology against IBM's hodge-podge of systems architectures and networks.

AT&T's great strength is that it has one of the largest bases of network management nodes in the world: PBXs in almost every U.S. corporation; central offices (COs) in every major telephone company; and an exceptionally powerful network backbone that reaches all around the world. In short, AT&T bestrides the telephony industry the way IBM casts its shadow over data processing. If it can translate such market power into network management sales, the company's future in this market will be stellar.

In private networks, we expect AT&T's share of underlying products such as PBXs, computers and personal computers, terminals, LANs, modems, multiplexers, and packet switches to remain stable. And we do not think the company has the sales strength to make its MIS/network management (NM) products sell. Less than spectacular growth is expected.

In public networks, AT&T already has a dominant share and is not expected to outperform the Network Systems Group, which sells COs and transmission products.

PRODUCT/SERVICE STRATEGY

AT&T's network management product strategy is built around its Unified Network Management Architecture (UNMA). UNMA is packaged in several modules called Element Management Systems (EMSs). The first of these, the so-called Network Management Protocol—which describes UNMA design criteria—has been introduced as having a broad range of products and services:

- ACCUMASTER Integrator, which provides a graphic view of the network and coordinates information from AT&T and non-AT&T systems (including SNA), reporting and isolating faults. The Integrator integrates all of AT&T's ACCUMASTER, STARKEEPER, Dataphone, and Acculink products, as well as Cincom's Net/Master and IBM's NetView. (Introduction scheduled for fourth-quarter 1989; entry cost about \$275,000.)
- STARKEEPER Network Management System manages DATAKIT Virtual Circuit Switch and ISN including fault, reconfiguration, and accounting management.
- Dataphone II System Controller for managing modems, multiplexers, and digital data sets in two-point, point-to-point, multipoint, and multiplexed networks.
- Dataphone II Acculink Network Manager for monitoring, controlling, testing, and reconfiguring backbone T1 multiplexers.
- ACCUMASTER Trouble Tracker, a trouble desk for tracking the performance of networks composed of a wide mix of PBXs, applications processors, and LANs.
- Centralized Management System offers facilities, traffic, terminal change, and cost management to users of AT&T PBXs, ISN, Audix voice messaging systems, and AP16 and 3B5AP applications processors.
- EPSCS Customer Network Control Center supports Enhanced Private Switched Communications Service customers configuration, fault, performance, and accounting management.
- Multifunction Operations System (MFOS) with OSS support for the following functions: network element maintenance; alarm processing; traffic data collection and reporting; congestion management; trouble tracking and administration; network element database administration; transmission facility maintenance; and billing data collection and reporting. MFOS supports a wide range of AT&T CO and transmission products as well as the Northern Telecom DMS-MTX, NEC NEAX 61, and Motorola EMX 250.
- Multifunction Operations Center provides configuration, performance, and fault management for the 5ESS.
- Accunet T1.5 Customer Controlled Reconfiguration mechanizes private network channel reconfiguration without service orders.
- Accunet T1.5 Information Manager (AIM) provides realtime network alarms and details status on Accunet T1.5 Service circuits (available fourth-quarter 1989);
- ACCUMASTER Consolidated Workstation, using the Integrator, simultaneously displays the management capabilities of six of the following systems:

AT&T Network Management Strategy

AT&T's Centralized Management System (CMS);
System 75 System Access Terminal;
Cincom's Net/Master;
Network management systems supporting Digital's VT100 terminal emulation;
Dataphone II System Controller;
Dataphone II Acculink Network Managers;
Dataphone II Dial Backup Unit;
ACCUMASTER Trouble Tracker;
STARKEEPER Network Management System;
Accunet T1.5 Customer Controlled Reconfiguration; and
Synchronous host computer access (3279 emulation), including IBM's NetView.

Release 2.0 will be available in May, 1989 for \$2,500; Release 1.0 customers can upgrade for \$1,000.

- MACSTAR I and MACSTAR II End Customer Management System (ECMS), a computer-based operations system for Centrex customers; MACSTAR I costs about \$60,000 and MACSTAR II costs about \$95,000.
- NetDirector End Customer Control System (ECCS) gives LEC customers the ability to monitor, reconfigure, and control their portion of the public switched network. It comprises three subsystems:
 - NetDirector Virtual Private Line Application (VPLA);
 - NetDirector Virtual Private Network Application (VPNA); and
 - 5ESS Switched Synchronous Data Application (SSDA).
- Digital Access and Cross Connect System (DACCS) IV Customer Network Controller, a central point for cross-connecting, rearranging, rerouting, and testing high-capacity lightwave signals.
- NetPartner Network Management System offers end users access to carrier Operations, Administration, and Maintenance (OA&M) using three 3B2/600 minicomputers, a DATAKIT VCS, and a Sun-3 series workstation. NetPartner supports AT&T

1ESS, 1AESS, and 5ESS switches as well as Northern Telecom DMS-100 switches. NetPartner can handle up to 40,000 digital and/or analog lines, 7,500 trunks, and 25 switches. NetPartner also interfaces with IBM's NetView network management system or its underpinnings, the Network Communication Control Facility/Network Problem Determination Application (NCCF/NPDA).

Before divestiture, products such as the Network Systems Group's (NSG's) DATAKIT Virtual Circuit Switch and the End User Organization's (EUO's) ISN were identical, but they have developed quite independently since. And, where AT&T is offering NetPartner to carriers, it is selling someone else's product, Cincom's Net/Master, to its corporate accounts.

Thus, although AT&T attempts a consistent and coherent product development platform in UNIX V, in network management the company has become a breeding ground for the heterodox and the dissimilar. UNMA is supposed to bring product development back under a single umbrella; it is also designed to allow AT&T to work effectively with other firms in areas where it is weak. AT&T has already negotiated several arrangements with other vendors; four of the most notable agreements are described below:

- Wang and AT&T have announced joint marketing arrangements for the Wang Integrated Image Systems using AT&T UNMA capabilities;
- QPSX Communications and AT&T have announced a joint venture under which AT&T will license QPSX's IEEE 802.6 standard 45M bps MAN;
- Cincom will market two essential components of AT&T's SNA Management Application for the ACCUMASTER Integrator. The first component, the UNMA application, was developed by Cincom specifically for the Integrator. The second component, Cincom's Automated Network Management (ANM), is actually a standalone component of Net/Master, which is Cincom's answer to NetView;
- AT&T acquired Paradyne to boost AT&T's offerings in T1 multiplexers, modems, statistical multiplexers, and channel extenders.

SALES STRATEGY

AT&T's sales strategy is driven by its belief that network management means different things to different people. To carriers, network management means cost reduction and improved "fill," or capacity utilization

AT&T Network Management Strategy

rates on the one hand, and the ability to offer new services to their customers on the other. To network users, by contrast, network management certainly means cost reduction, but more importantly it means the improved operational capability to meet the business plan.

In effect, AT&T's Network Systems Group (NSG), which sells to carriers, can sell the "comfort level" that carriers need to buy complex network management systems and services. But AT&T's End User Organization (EUO), which sells to just about everybody else except the federal government, is having a much harder time convincing its customers to make "the right choice" in network management.

Also, while AT&T says it is selling a unified network (UNMA), it does not, in fact, have a unified network management strategy. Indeed, each AT&T line of business has its own agenda, is organized differently, has its own lab, and has its own ideas about customer relations.

The result: AT&T has two sales strategies in network management.

In the mature telco market, NSG dominates and will use UNMA to defend its base from the competition. It will also use UNMA to prevent troublesome competitors such as Northern Telecom from using their large base in switches to migrate into important value-added markets such as Operations Systems, a \$500 million plus business for AT&T.

In the growing corporate market, by contrast, EUO is using UNMA to break IBM's powerful grip on MIS.

The contrast between the two lines of business could not be more marked. NSG is focused, well run, and profitable. The carriers, its target market, are relatively homogeneous, and the company has a well-established and effective support infrastructure to serve them.

EUO, by contrast, lacks sales focus, has a bloated management (the EUO Chief Financial Officer's department alone contains 12,400 people), and loses money just about everywhere except long-distance services.

EUO constantly shifts from one target market to another. The Modified Final Judgement, which broke up the Bell System in 1982, initially forced AT&T to operate two end-user sales forces: one, called Information Systems (IS), sold hardware and software; and another, called Communications, sold network services.

The IS people spent most of their time chasing down new business, ignoring their large base of rental customers. The Comm people, by contrast, spent most of their time trying to increase the network usage of existing customers. With each passing year, the two groups grew further and further apart.

AT&T was allowed to merge the two in 1986, a process that created almost as many problems as it solved. By virtue of its much larger proportion of company revenues, the Comm people now dominate EUO sales strategy. As a result, the EUO today is more concerned about protecting its base than with going after new business. Most growth opportunities, therefore, go by default to IBM, Northern Telecom, MCI, and Sprint, among others.

AT&T has also been under attack for its pricing: regulation makes pricing flexibility difficult and usually forces the company to publish details others would keep confidential. The firm has been forced to publish its winning bid for the federal government's FTS 2000 network, for instance, something which may expose it to litigation by its competitors. But the flexible pricing policies in its Tariff 12 have been allowed by the FCC.

Today, the EUO sales force is 22,000 strong and in 292 locations nationwide. It is supported internally by 12 data centers with 160 mainframes, 1,060 3B computers, 11,000 personal computers, 9,000 terminals, 209 DATAKIT/ISN LANs, a 40-node wide area network (WAN), and 500 dedicated circuits.

Yet another reorganization is in progress, however, and the impact of this on the sales force is not clear. The EUO has been broken into 15 line divisions in charge of fourth-level management. Bob Allen, the AT&T CEO, has thus been able to eliminate upper management clutter and give younger, more aggressive line chiefs a clear mandate, real bottom-line responsibility, and a timetable for returning their operations to profitability.

The new AT&T, therefore, should look a lot like IBM: several product-oriented lines feeding one sales organization. But we do not think matters will stay like this for long. Some product lines, such as PBX, are chronic money losers and may not be susceptible to turnaround without radical restructuring and massive layoffs. Even so, profitability may not be possible within Allen's timetable, and these divisions may be divested. Indeed, we expect several such sell-offs within the next 24 to 36 months and are calling this "The Second Divestiture" of AT&T.

If enough product groups are sold, probably to European and Japanese companies looking for good op-

AT&T Network Management Strategy

portunities in North American telecom markets, sales operations will have to be restructured. We expect that much of AT&T's sales positioning will shift from hardware sales to network-based services and to a much broader use of OEM'ed products and joint ventures planned, like the Wang and Cincom ventures, under the UNMA umbrella.

Strengths

The AT&T-Communications account base: almost every company in the U.S. uses AT&T services in one form or another. Few firms in or out of telecommunications have as much reach.

The AT&T-NOG network backbone: the network gives AT&T an exceptionally powerful platform from which to leverage its unequaled account base.

The Network Systems Group: a well-managed, profitable, and highly competitive operation that maintains high standards of customer service.

UNIX: for all of AT&T's heterodox implementations of it, the company can do all its network management development work under the UNIX V operating system umbrella. IBM, by contrast, must work in dozens of operating environments and has only just begun to define intersystem communications and applications development protocols.

OSI: because AT&T's technology platform is fairly uniform, and because its OSI goals are clear, the company is well placed to make network management acquisitions—such as Paradyne—or enter into technology joint ventures, such as the ones with Sun and Cincom. This enables AT&T to fill holes in its network management product line with relative ease.

Weaknesses

Constant EUO reorganization: AT&T people are simply not used to the company's current pace of reorganization, and many are sure to get lost in the shuffle. Customer service levels will suffer as a result.

The threat of further divestiture: this is certain to make AT&T-EUO staffers worry about their futures with the firm. Again, customer service levels are bound to suffer.

Account management: in spite of the enormous breadth of its EUO accounts, AT&T has very little

depth in each, especially when compared to arch competitor IBM. Some of AT&T's best network management ideas, such as its 6500 front-end processor, do poorly for this reason alone.

Sales: weaknesses are evident throughout the EUO. Sales of everything from key systems to PBXs and computers are lackluster, and profits are nonexistent.

Sales force incentive programs: salespeople have not been offered incentives, or even given quotas, for ACCUMASTER products, the top of AT&T's network management line. Said an AT&T product manager, "They don't have to be; these products will sell themselves."

Limited management-customer interaction: EUO ACCUMASTER Integrator product management did not visit any customer operations until after the ACCUMASTER Integrator was introduced.

Poor marketing support: EUO sales representatives report that headquarters marketing support is useless despite (or perhaps because of) its voluminous staff. Specifically, a series of internally generated newsletters on vertical markets contain no information pertinent to sales. Expensive resources are wasted.

High overhead costs: someone has to pay for AT&T's vast staff organization. The company must support 35,000 nonrevenue producers, of which 12,000 alone work for one line of business, the EUO.

FUTURE DIRECTIONS

While IBM has superior account management skills and better penetration of large accounts, AT&T's uniform operating system platform allows it to demonstrate superior customer features and benefits. To sell network management systems effectively, however, AT&T must shorten the lines of communication between its customers and its labs and between its customers and its upper management. Even though AT&T's technology platforms allow for strong internal development and interesting joint ventures in network management, apart from NSG, AT&T's sales organizations are not well positioned to take much advantage of the opportunity.

Thus, AT&T may not actually be capable of selling its UNMA products and services for much more than simple call detail recording (CDR) and call queue management activities. □

BellSouth Network Management Strategy

This report will help you to:

- Assess the effectiveness of BellSouth's initiatives to establish a uniform and powerful network.
- Discover how BellSouth plans to wrestle with two opposing forces—its regulators' goals and its customers' needs.
- Compare BellSouth's position to that of other LECs.

MARKET POSITION

Like all Local Exchange Carriers (LECs), BellSouth is being torn apart: The goals of its regulators—universal, basic service—and the needs of its customers—maximum flexibility in price, packaging, services, and coverage—are moving farther and farther apart. Thus, efficiently managing its own network *and* those of its customers is becoming impossible.

To accommodate its two very different masters, the company has had to rewrite its mission and redesign its network program from top to bottom. BellSouth has completely rethought its operating philosophy from the chairperson to the line maintainer. Network management is at the center of that new line of thought.

This report examines BellSouth's new philosophy and assesses its strengths and weaknesses. The conclusions assume that BellSouth's network management revenues will grow fairly evenly over the next five years, with most of that growth in private branch exchange (PBX) and key telephone systems (KTSS). Also, the analysis assumes that the Modified Final Judgement (MFJ) is not modified to allow simpler Regional Bell Holding Company (RBHC) participation in private network management markets, and that the Public Utility Commissions (PUCs) continue to object to network modernization on the grounds of "overbuilding."

By its own admission, BellSouth in its early years suffered from an unhealthy combination of complacency and misdirection—a condition that pervaded all carriers spun off from AT&T.

While the firm saw, and continues to see, huge opportunities for adding value to its customer services as the business of networking grows and changes, the thinking of its regulators has remained unchanged for decades. Where the regulators' and customers' needs once coincided, little, if any, congruence exists today.

Thus, it is almost impossible for BellSouth to have a useful direction, let alone achievable goals. The company cannot be too customer oriented: This would offend its regulators. And it cannot be too regulation oriented: This would offend the customers. No strategy that attempts to bridge this gap is sound; its inherent flaws make failure inevitable.




Figure 1. BellSouth network management revenues, 1988 to 1993 (\$ thousands).

BellSouth Network Management Strategy

This problem was not obvious during the years immediately following divestiture, when the firm performed well almost in spite of itself and became complacent. During 1987, however, management decided that to avoid a major crisis, it must revise its strategy.

Since it cannot adequately serve two masters, BellSouth decided to redefine the problem, as it were, by concentrating wholly on its cash flows. In its simplest sense, BellSouth has decided to become financially driven. This has meant abandoning a lot of high-flown ideas about new network management services for its customers; the wholesale revision of network upgrading plans (Southern Bell canceled plans to replace 192 AT&T 1AESSs with digital products, for example); and ruthless cost reductions and layoffs.

While the company still pays public homage to the customer with slogans such as "To be the best there is, the customer decides," BellSouth's number one concern today is work force reduction and cost control. The restated mission boils down to providing "flexible, economical, and reliable services to our customers." Services come first; customers come last.

To meet its new objectives, BellSouth is launching new initiatives for managing its own network. It is automating the network as much as possible, to eliminate overstaffing and to improve capacity utilization or "fill" rates. Here the firm is on strong ground: It owns and controls its network and has a high degree of flexibility in deciding how to improve network efficiency.

Among its customers, however, BellSouth is in a rather different position. Regulation simply doesn't allow BellSouth the luxury of being too concerned with its customers' needs. It has a large number of customers with data transport requirements and has a broad base of data transport services for them, but can add little value otherwise.

To do what it can, BellSouth is forming strategic alliances through the unregulated side of its business. With Digital Equipment Corporation, for example, BellSouth is offering network management packages for DECnet customers. Also, BellSouth is designing and selling a fiber network management software system in conjunction with RAYNET. BellSouth is also trying to offer low-cost database access in a joint effort with Telenet.

South Central Bell has about 3,800 Digital Data Service (DDS) lines and 5,400 T1 lines in service; Southern Bell has 5,400 DDS lines and 1,500 T1 lines. Revenues for both DDS and T1 are \$55 million annually. In addition, 15 percent of BellSouth voice grade private line revenues (another \$36 million worth) are derived from Dataserv customers using voice grade for slow speed data (usually 9.6K bps or less). Thus, BellSouth data transport revenues are running at about \$91 million annually—not much for a firm with annual revenues of \$13.7 billion.

The company clearly sees the opportunities in the data communications market and wants to leverage this base. But it is difficult to see how this can be done. The firm offers Electronic Tandem Networking (ETN) services now and, once Signaling System 7 (SS7) is in place, will offer Private Virtual Network Service. BellSouth has a handful of Central Office-Local Area Network (CO-LAN) customers as well. But BellSouth's SS7 plans have been scuttled with its 1AESS replacements, and the alternatives are limited.

BellSouth has already done just about the only thing it is currently able to do for its customers—improve customer service levels and cut costs by integrating its regulated services sales force with its unregulated services sales force.

For itself, BellSouth's main goal seems to be maximized throughput and maximized cash flow. To do this, BellSouth is working on two fronts: the implementation of SS7 and a new Operations Support System (OSS) environment.

The OSS is called the Signalling Engineering and Administration System (SEAS), and it will primarily be used by BellSouth's Signalling Engineering and Administration Center. The SEAC is an operations center overseeing the CCS 7 network and monitoring all CCS




Figure 2. BellSouth access line growth 1984 to 1988 (thousands of lines added annually).

BellSouth Network Management Strategy

7 network performance. As offices are cut over, SEAC monitors the network and load distribution on SCPs and plans reconfiguration where required.

Today, BellSouth's CCS 7 network consists of two Regional STPs (RSTPs) and two Local STPs (LSTPs), one of each being located in Atlanta, Georgia and Birmingham, Alabama, which serve Service Switching Points (SSPs) in Chattanooga, Tennessee and Atlanta and run under the SEAS located in Charlotte, North Carolina. SEAS is monitored from the SEAC in Atlanta.

The first CCS 7 application is 800 service, for which BellSouth has an 800 Service Management System located in Kansas City, Missouri.

Short of ISDN, however, the whole network is run on a "virtual" ISDN in which traffic is sent over 56K bps links and signaling traffic over 9.6K bps "special service" links. This internal network is being leveraged for BellSouth customers through the Simultaneous Digital Voice and Data (SDVD) "pre-ISDN" virtual 2B+D service, which the carrier is offering to improve basic voice and data transport services to large customers. To make the interLATA portion of this service work, BellSouth is working with US Sprint, among others.

BellSouth has targeted SDVD network management services to:

- large commercial customers;
- small business and residential customers;
- universities and colleges;
- local and federal governments; and
- public utilities.

SDVD is meant to bridge the gap between packet switched services, such as Pulselink; various DDS and analog circuit switched services; and ISDN-based network management offerings which may be allowed in the future.

But SDVD may have a longer life than planned. BellSouth was to have increased depreciation allowances on its COs, enabling it to replace existing analog electronic COs, such as the 1AESS with what it calls "fourth generation" systems in the 1989-1990 time frame. Presumably this meant buying SuperCore-based DMS-100s from Northern Telecom and 5E5 Generic 5ESSs from AT&T. This program came to a sudden halt, however, when regulators accused BellSouth of overbuilding its network and refused it per-




Figure 3. BellSouth network management expenditures, 1984 to 1993 (\$ millions).

mission to make the planned changes. BellSouth's network management offerings may just have been KO'ed.

PRODUCT/SERVICE STRATEGY

BellSouth's overall strategy in network management will initially be all infrastructural: The telco wants to improve its asset management and improve fill rates.

By the end of 1989, BellSouth's entire network will be programmable: all its electromechanical signaling points will be replaced.

BellSouth plans to place more fiber in the trunk and feeder portions of the loop; the company is aggressively pushing fiber in the distribution portion of the loop as well. Indeed, BellSouth data shows fiber proving in over copper in the distribution portion of the loop as early as 1990.

Loop, trunk, and switch operations will be integrated with planning, engineering, and operations systems disciplines. Dozens of separately managed OS systems, such as those listed below, will be better integrated through SEAS. As a result, the percentage of remotely reconfigurable devices on the network will increase from about 1.6 percent currently to some 20.0 percent in 1990:

- Computer Systems for Mainframe Operations Support (COSMOS);

BellSouth Network Management Strategy

- Computer Access System (CAS);
- Direct Order Entry (DOE);
- Installation Maintenance Center/Operations System (IMC/OSS);
- Line Information Database (LIDB); and
- Trunk Integrated Record Keeping (TIRKS).

STRENGTHS

Management is BellSouth's greatest strength. It has the will and the ability to make tough decisions, such as deep staff cutbacks, that almost always indicate toughmindedness and a clear sense of direction.

BellSouth's clear commitment to a uniform and powerful network is also an important asset. This may sound trite—everyone is committed in this way—but BellSouth is closer to having a 100 percent stored program control (SPC) network than any other Regional Bell Holding Company (RBHC).

Good experience at managing a "virtual" out-of-band network internally.

On-line connection to Northern Telecom's First Application System Test bed in Raleigh, where BellSouth can test proprietary network management applications.

WEAKNESSES

Regulation is the main force that holds BellSouth in check. The carrier simply cannot do much more for its data customers than offer cleaner channels, or more of them, at better prices.

Experience shows that network bifurcation, and the resulting bifurcation of sales and marketing teams, has a negative effect on sales. Wherever the customer requires any hardware at all (RBHCs cannot even touch inside wiring by themselves), sales and service turn into administrative nightmares—exactly the opposite of what network management is supposed to accomplish.

FUTURE DIRECTIONS

BellSouth is in a tough, but interesting, position. It has demonstrated the will and the management ability to get a firm grip on its own network. But it is straining at the leash of regulation. And with Judge Greene saying firmly that telcos cannot even offer interLATA information gateways, there is a limit to what it can deliver apart from basic transport services. □

IBM Network Management Strategy

This report will help you to:

- Discover why IBM is aggressively entering into the network management market.
- Evaluate the effects of IBM's repositioning and product marketing strategies on your future purchase decisions.
- Evaluate the present and future risks incurred by IBM and its customers as the company changes direction.

IBM is in the midst of one of the most dramatic repositionings in the history of the information industry. The computer giant has entered the network management field and plans to become the world's premier value-added supplier of network and information management services. IBM's entry into the network management arena will affect competitors' strategies as they assess the significance of the move on their markets.

MARKET POSITION

IBM is aggressively entering into network management because the company cannot afford to base its continued financial health on the slow-growing mainframe market. Aware that computer hardware and software constitute only a fraction of the value of a fully functioning network, IBM sees growth and profit opportunities in network management.

Current IBM customers will see a dramatic change. To position itself better to serve customer network needs, IBM has already begun to offer competing products and will soon offer many more. Systems integration will take precedence over IBM's own products in an effort to offer customers maximum network leverage.

For companies now competing with IBM, the changes will be equally dramatic: they may soon find in IBM their biggest sales channel. Others, such as the tele-

phone companies that seek the same markets, will soon feel the full weight of IBM's competitive power.

IBM leads the world's mainframe market, the midsize minicomputer computer market, and the personal computer market. IBM reached its lofty position through many skills, notably unsurpassed account management. When IBM decides to bring its prodigious resources to bear on new opportunities, the entire industry trembles.

Orchestrating those changes, however, can be tricky. The company does not want to move so fast that it

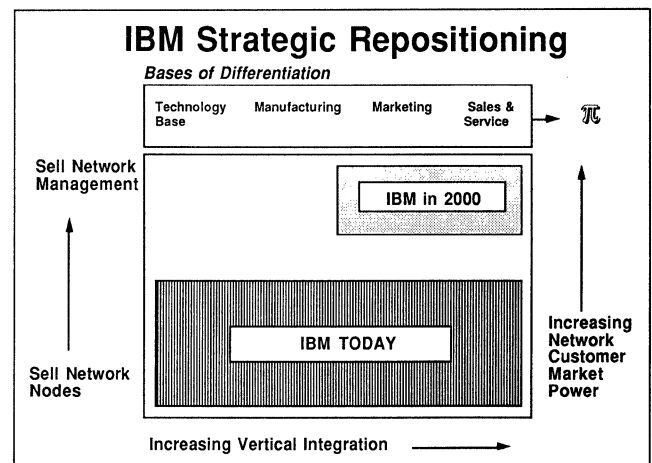


Figure 1. IBM's strategic repositioning.

IBM Network Management Strategy

hurts its core business and the principal source of its cash flow—large computer systems. Nor does IBM want to move so slowly that it becomes vulnerable to countermoves by major competitors, especially Digital Equipment, EDS, Sun, and AT&T.

At the same time, the market for information management products has changed. What was once IBM's greatest strength—its endless market segmentation and the extreme applications orientation of all its products—has become the firm's greatest weakness.

Today, MIS managers are under enormous pressure to reduce their operating costs and to increase the return on MIS investments by using network management tools to optimize their complement of computing systems. Since IBM does not have a homogenous product base, optimizing its customers' MIS operations while concurrently engineering its own repositioning creates high risks.

And if IBM faces risks, so do the customers it serves. This report identifies those risks and isolates critical network action points.

IBM'S PRODUCT/SERVICE STRATEGY

IBM's corporate mission can be simply stated: maximize network throughput and skim off the value added.

To achieve that goal, product flexibility and ease of network management are crucial. IBM's product/service migration strategy must shift from selling more MIPS for less to selling more megabits per second (Mbps) for less. Thus, an enormous portfolio of products that were not designed to work together presents a huge liability.

Some IBM systems, such as mainframes, are "top heavy": they control from the top of the network, as it were, leaving little room for distributing control to the terminal base below. Indeed, IBM's mainframe terminal base of 3270s was intentionally limited in capability.

Other IBM systems, such as PCs, are "bottom heavy": all control resides in the terminal. Indeed, ceding PC network control to an IBM mainframe often means abandoning PC functionality—in short, wasting a lot of money.

As it is, IBM systems are top heavy, bottom heavy, and everything in between. It also has disparate customer bases to answer to, such as the millions who have spent a little to buy its smallest system, the PC; the thirty thousand or so who have invested small fortunes in

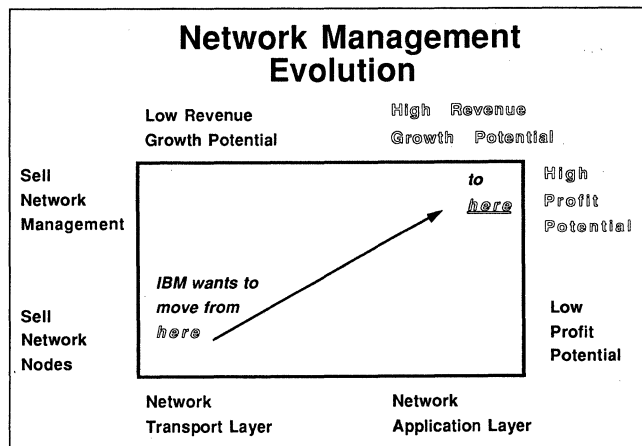


Figure 2. Network management evolution.

System/370-based SNA networks; and the several hundreds of thousands who have bought the mid-sized systems that span the gap between the extremes.

To fulfill its goals and build a system that can successfully interconnect all of its products and offer the right degree of interoperability to satisfy all its customers, IBM must do one of two things:

- redesign everything—a complex task that makes all existing products obsolete; or
- modify every product to accept a common overlay, which means a lot of modification for some machines and a little modification for others—an inefficient and expensive alternative, but one that will not render everything obsolete all at once.

The first option might have been acceptable even a decade ago. In the late sixties, IBM dropped the System/360, and in the late seventies, did the same with the System/34. IBM succeeded in those strategies because it could, at that time, control customers' buying decisions. When IBM moved, its customers moved with it, even when, as with the introductions of the /360 and /370, those moves caused heavy financial burdens. Now, running parallel with the increasing complexity of networks, the increasing market power of the customer loosens much of the hold IBM fastened on its customers' buying decisions.

Competition and the diffusion of processing demands into smaller systems have impelled many of IBM's largest and most profitable accounts to resist ordering upgrades to hardware and software on command. Simply put: customer market power is gaining the upper hand.

IBM, therefore, has had to adopt the second option by default. Doing something by default, however, does not impress an increasingly demanding customer base,

IBM Network Management Strategy

which is almost certain to reject machine modification almost entirely in the long term. There are limits to how far and how fast IBM can move on its own.

Thus, IBM's sales strategy has shifted from an extremely proprietary market/applications to a product/technology orientation, which includes a heavier and growing emphasis on the products and services of others. This shift will revolutionize customer sales and service relationships with its customers and will offer network managers who rely heavily on IBM mainframes—indeed on any IBM products—much-needed flexibility and new reasons to shop IBM.

There will be a price, however. Before IBM can complete the conversion of its products, customers will be asked to pay for a vast series of upgrades and enhancements just to make the IBM network functional. Vendors of competing architectures, such as Digital Equipment with its Enterprise Management Architecture to be launched later this year, announce that they incur no similar overheads. Many of IBM's customers are certain to agree and will decide that the cost of staying with IBM is just too high.

The difference for most, however, will be the enormous cost of replacing their IBM systems and the corporate operations and procedures that have grown up around them.

For the undecided, we expect that the influencing factor will be IBM's long tradition of excellent service. These companies would rather keep good service, pay IBM's price for network optimization, and wait.

In summary: IBM's entire product and service strategy is network management. To put this strategy in place, however, IBM has to bet its entire customer base and pray that those customers will stick with the company during a difficult period and that the competition is unable or unwilling to take advantage of IBM's position.




Figure 3. IBM network management revenues, 1984-1993 (\$ billions).

IBM'S NETWORK MANAGEMENT PRODUCTS

For a company that places such importance on network management, IBM's list of network management products is surprisingly short, consisting of two network alert and monitoring software packages and a T1 OEM'ed from N.E.T.:

- NetView, a host-based system;
- NetView/PC, a PC-based system for managing distributed networks; and
- 9753 IDNX T1 Resource Manager.

This list, however, is deceptive. The company has spent the better part of the last five years—since September 1985, when the first major thrust into network management was proclaimed—announcing a vast array of product and service upgrades. Major announcements have occurred almost monthly since 1985 and are often published in batches of 500 or more, affecting virtually every item in IBM's inventory. Indeed, Systems Applications Architecture (SAA), announced in March 1987, can be viewed as a major network management vehicle driving the interoperability of all IBM's systems.

NetView is an enhancement of VTAM (Virtual Telecommunications Access Method), the base for the major IBM communications subsystems, and runs under all of IBM's extant host operating systems, including VSE, VM, and MVS. NetView consists of several subsystems:

- NetView File Transfer Program
- NetView Distribution Manager
- NetView Performance Monitor
- NetView Access Services
- NetView Network Definer
- Network Design and Analysis (NETDA)
- TSO/CMS Information Management
- Voice Accounting Application
- Host Command Facility

Running under VTAM, NetView is designed to manage the following current and discontinued products:

- System/370 and System 9370 mainframes

IBM Network Management Strategy

- 3174, 3274
- AS/400 and System/36 and -/38 departmental computers
- 43XX midrange computers
- 3710 and 3708 cluster controllers
- 4700
- System/88 fault-tolerant async host
- Series/1
- System 8100
- IBM LAN Manager
- IBM Transmission Network Manager
- OS/2

In essence, NetView manages synchronous SNA networks; NetView PC manages everything else.

NetView/PC contains the following elements: ROLM Alert Monitor, ROLM Call Detail Monitor, IDNX Alert Monitor, and Series/1 Remote Manager.

Fundamentally, NetView/PC offers improved CBX functionality and LAN and T1 management and also interfaces with competing systems and networks. Between NetView, a monitoring and alarm system, and the 973X IDNX T1, which cuts leased line costs, IBM has moved to answer customers' two most important demands—improved management information and leased-line cost reduction.

Any effort to internetwork all IBM's products, no matter how sophisticated, will inevitably be clumsy and expensive. NetView alone can consume 10 to 15 percent of a 3090's processing power, and accessing NetView PC can take up as much as 400K bytes of a PC's RAM, a great deal in an environment in which PCs rarely carry more than 1M byte and where many operating system/application packages will take 800K to 900K bytes, and data files all the rest.

As another example, recent improvements in the SNA Network Control Program (NCP) to replace PU2.0 functionality with PU2.1 free up the mainframe and add flexibility to the network but extract 200K bytes of RAM overhead from every PC attached.

IBM Information Network: Initially, IBM had some difficulty positioning I-Net for maximum effectiveness. A few years ago, IBM tried to follow its move into

mass market interexchange services through SBS by turning I-Net into a mass market VAN in a proposed joint venture with British Telecom. The British government and most of the British press objected vehemently, as did, predictably, IBM's competitors. IBM retrenched, turning I-Net into a primary vehicle for SNA-based VAN and network management services for *Fortune* 500 "enterprise" accounts.

The I-Net backbone comprises 36 S/370 hosts (3080s and 3090s) connected through 56K bps and T1 leased lines. This network connects 4,500 U.S. organizations and serves 320,000 users, including a large number of IBM's own staff. Worldwide, full SNA service is offered in 21 countries, and packet (X.25) connections are offered to 55 others. Over 500 types of terminals are connected, and dial-up access is available at 300, 4800, and 9600 baud.

IBM's Network Management System Concept: IBM has grouped its customer network management requirements into six major "disciplines" or applications to be served by its NetView service umbrella. These management functions include operations, configuration, problem, change, performance, and accounting.

To understand IBM's network management product/service positioning, it is necessary to examine each of these network management applications separately.

Operations Management: a network alert function performed by suppressing unneeded messages or automating message responses by using the NetView application interface. IBM performs this function by using the Host Command facility, NetView Access Services, SNA Applications Monitor, and the NetView Inter-System Control Facility to provide remote, unattended operations; hardware monitoring and control; generic service point commands; and network application access and control.

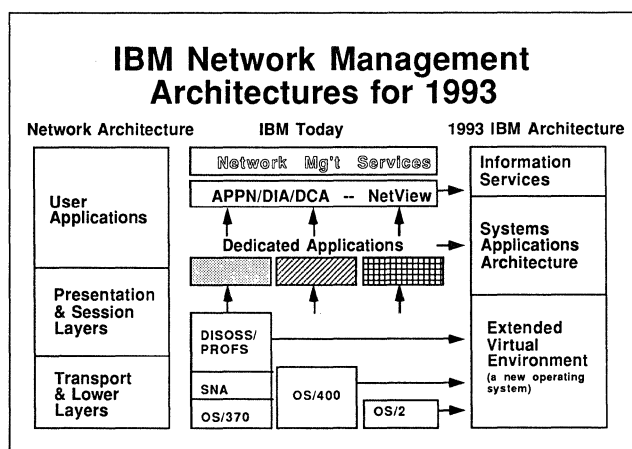


Figure 4. IBM network management architectures for 1993.

IBM Network Management Strategy

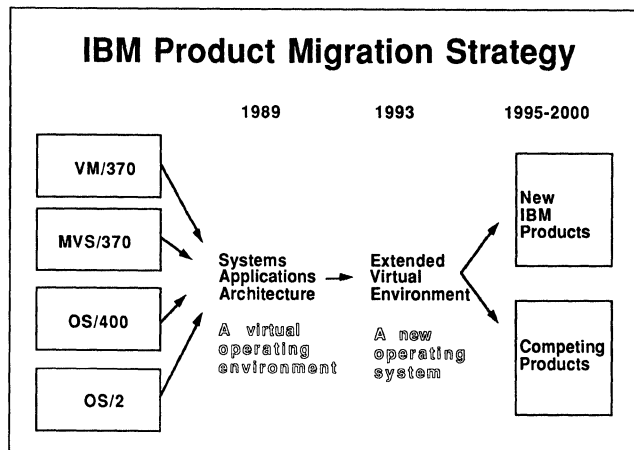


Figure 5. IBM Product Migration Strategy.

Configuration Management: the capability of the network to describe itself and its constituent parts. IBM performs this function by using host CICS, TSO/CMS IM, NETDA, and the NetView Session Monitor to provide dynamic session configuration; modified network system definitions; and display of physical and logical network connections.

Problem Management: primarily a statistical function that allows the network to detect, track, and correct problems. IBM performs this application by using the host Information Management System running under TSO/CMS, the host Service Level Reporter, and the NetView Performance Monitor to provide problem alerting; automated recovery; problem diagnosis; problem tracking; action recommendations; performance analysis; generic service point commands; and trend analysis.

Change Management: for software updates and customization data. IBM performs this function by using TSO/CMS IM, the NetView Distribution Manager, and the NetView File Transfer program to offer network plan definition; automatic retry; security; bulk file transfer; batch submission; database update and refresh; and change tracking.

Performance and Accounting Management: records response time data and resource utilization statistics, including response time measurement; component utilization; work load measurement; availability measurement; batch throughput measurement; system tuning; capacity planning; accounting; and service-level management.

IBM performs these functions by using the host Network Design and Analysis capability, the NetView Performance Monitor, host Service Level Reporter, and the NetView Session Monitor to offer performance

monitoring and analysis; network accounting; voice facilities accounting; and performance planning (VM OS only).

“Disciplines” apart, IBM’s customers really want a cheaper, more reliable network. To deal with these issues, IBM has put a major effort into redesigning the heart of its network: the SNA relationship between terminals and hosts through the 37X5 front-end processor group.

Little in the way of leased-line savings can be accomplished so long as the front-end processor could not handle multiple hosts; redesign the network in the event of nodal failure; and operate, in effect, in a so-called realtime, reliable mode. IBM has performed major upgrades on the 37X5s and on cluster controllers in recent years, as well as the way in which the network runs under the host VTAM program.

Today, for example, the newer releases of the Network Control Program do not require the entire network to be taken down simply to redesign or to alter terminal application sets, and front-end processors can address multiple hosts. This may seem a fairly trivial exercise, and one that should have been accomplished in the early years of SNA, but it is a measure of IBM’s technical problems that customers have only recently been relieved of these costly burdens.

As an added benefit, the repartitioning of the network management load under the latest version of NCP also puts more of this load onto the 37X5 and the terminal and reduces dependence on host-resident VTAM. This procedure, in turn, can reduce leased-line costs and mainframe overhead.

SALES STRATEGY

IBM has completely redesigned its worldwide sales operations several times in recent years to better position the firm to benefit from opportunities in network management.

In the latest incarnation, the newly created U.S. Marketing and Sales Group, in charge of all sales in the U.S., focuses completely on IBM’s best market—its large accounts. All small accounts are now being encouraged to work through third parties.

IBM’s major account focus allows account sales teams to start by assessing each company’s organizational structure and strategy. The team then tries to position the customer to accept a “network philosophy,” as it were, that is compatible with the organization or that might even help to drive it. IBM then uses the network

IBM Network Management Strategy

philosophy as a framework to sell everything from mainframes and direct access storage devices (DASDs) to wire, cable, and 3270s.

This approach is a big change from when IBM's sales forces focused on applications or sets of needs and sold to those needs in specific user groups in each large account. IBM now understands that since the network is the primary vehicle for services, and since coherence is essential to the network, its systems must serve the network; the network cannot serve its systems. Sales forces have been restructured accordingly.

Also for the first time, IBM sales, marketing, and product development people are all marching to the same tune and playing that tune for IBM customers. Thus, IBM customers are getting a clearer, more committed message than they have heard in years.

This sales consistency has been combined with a new openness. Not only do IBM salespeople routinely recommend non-IBM products, but they actually sell them, as in the case of N.E.T.'s IDNX T1 and TSB's Call Collector 3.

IBM salespeople now discuss future product direction. Such frankness is essential to selling into networks that demand backward compatibility and commitments to the future, but it was forbidden at IBM not so long ago.

STRENGTHS

In recent years, IBM has shown its toughness in a variety of ways, all designed to give the message to its customers and to its own people that it is a firm with a mission and will not be pushed from its chosen paths. IBM has:

- reorganized from top to bottom, infusing the lines with a new centrist discipline designed to pilot corporate repositioning.
- cut large staff organizations and *increased* its already prodigious sales force by approximately 20,000 former middle managers.
- dropped a range of products not central to the core of its network management mission.
- completely upgraded most of its remaining products.
- restructured its entire sales force into new account management teams.
- pushed sales of all small-ticket items out to third parties to give its sales force even more focus.

- divested interests in companies peripheral to its network management vision, such as SBS, Intel, and Rolm.
- acquired companies in areas of vital importance, such as PacTel's Spectrum Services network monitoring capability.

WEAKNESSES

The major weaknesses that IBM must address are as follows:

- Decades of endless market segmentation have left the firm with dozens of incompatible architectures.
- A series of missteps in telecommunications hardware and services—from SBS through to Mitel and Rolm—have left the company with a stained record in network management.
- Critical limitations in SNA have required much work and may never be surmounted without replacing the System/370 architecture.

FUTURE DIRECTIONS

From a purely technical aspect, it is hard to imagine a company in more trouble. Compelled to shoehorn its own product design failings into the market, hoping that its customers' investment in them are so great that they will be reluctant to move while IBM sorts itself out, IBM is easily the weakest player in the network management business.

But, unlike most others, IBM knows what is wrong and is going after solutions with a vengeance. Indeed, few players have demonstrated such persistence in the face of difficulties monstrous enough to handily destroy a company with less management strength.

IBM's customers, therefore, have several issues to review before deciding just how far they will go with IBM into network management. These issues are:

- IBM may not get all its products to work together before the beginning of the next century, which may not fit its customers' timetables.
- If IBM does not get its products fully internetworked by, say, 1995, the overhead costs may be so high that customers might decide to look elsewhere for their network management products and services.
- Management doesn't live forever. If IBM's management changes before the repositioning is complete

IBM Network Management Strategy

and new management is not as committed, or if the repositioning fails and management changes as a result, what effect will this have on IBM customers?

NBI's view is that IBM's repositioning will be completed by the end of this century. See Figure 4. Between now and 1995, and to a lesser extent between 1995 and 2000, IBM will completely replace all the products running under its SAA overlay. See Figure 5.

This replacement will occur in three phases.

During the first phase, which has already begun with the launching of the 9370, AS/400, NetView, and OS/2, the power of various products will be increased, sometimes unevenly, in an effort to make everything powerful enough to run under SAA. OS/2, for example, an essential SAA platform for IBM's PCs, will require an unheard of 4M bytes of RAM just to boot.

Other products, such as the AS/400, will be upgraded to straddle a difficult line between backward compatibility with preceding products, such as the Systems/36 and -/38, and future product compatibility under SAA.

In this phase, IBM is establishing SNA as a base for complex networking and network services; integrating OSI for interconnection with non-IBM systems; supporting TCP/IP to allow TCP/IP users full access to IBM networks; providing protocol conversion and custom gateways for selected customer networks; and using open architectures and OSI standards to integrate network management for IBM and non-IBM systems.

During the second phase, soon to be launched and probably scheduled for completion within the next 36 months, IBM will use the Common User Access (CUA) protocols of SAA to uncouple as much of its software from underlying hardware as possible. NBI predicted the launching of this phase in 1986, calling it the Virtual Environment. Today it is called SAA.

Under CUA, IBM will replace systems-based applications development with what it calls "case-based applications development," designed to run on any SAA-compatible system, especially on systems running under four of IBM's current operating systems—VM, MVS, OS/2, and OS/400—and to be completely workstation based.

The CUA will break up the systems/applications orientation of IBM's products and allow the creation of a homogeneous operating environment that works

across product lines. We call this stage the Extended Virtual Environment, or XVE.

Once CUA has uncoupled IBM's products from their applications sets and operating systems, IBM will begin the third or XVE phase of its product migration.

Beginning 36 months from now and lasting through 1995, the company will "harden" its CUA-based virtual operating environment, dropping the four operating environments previously listed for the single operating system, XVE. This approach will mean replacing every product in its portfolio with something that better exploits XVE.

During this third phase, IBM will also push as many of its current competitors as possible to adopt the CUA, allowing them simple access to all IBM networks. XVE will be designed to facilitate that access. In effect, like UNIX, the XVE operating system will become IBM's network management framework, and IBM's network management and operating environments will become indistinguishable.

By 1995, IBM's products and services will look more like those of Digital Equipment or AT&T today. Its success in this transformation, and the success of its repositioning in network management, hinges on the workability of XVE.

Indeed, shifting the product development focus away from applications and toward the interoperability of its systems means perfecting network management. Without it, neither IBM's products nor its customers can be profitably migrated.

Whether the company completes all phases is vitally important for network managers. The first phase will allow IBM's customers to retain their existing investments in applications software, largely because the underlying architectures and operating systems will not change much. During the second phase, IBM will prepare to change-out its entire product line. During the third phase, this change-out will be completed.

IBM's customers will have to examine their options carefully. Many will decide early to minimize their exposure to IBM systems and will shift to other network management suppliers while IBM's offerings remain relatively unsophisticated.

NBI expects AT&T with its ACCUMASTER product line to be one of the main beneficiaries. ACCUMASTER and the ACCUMASTER Integrator use the established networking power of UNIX to offer a broad array of alert monitoring and can manage async and sync, voice and data, networks. AT&T has the network management skills that IBM is just be-

IBM Network Management Strategy

ginning to acquire. But AT&T has nothing like IBM's account management skills to take full advantage. The two companies are on a collision course: both want to seize the network management high ground. The fight will be even—AT&T has the technology;

IBM has the accounts. MIS management faces a hard choice between buying excellent, powerful products (from AT&T) or buying service levels that are probably even more excellent and powerful (from IBM). □

MCI Network Management Strategy

This report will help you to:

- Compare MCI's network management strategy with those of other major vendors.
- Understand MCI's NMS pricing strategy and how the company plans to grow in the network systems market.
- Evaluate MCI's network management products and services.

MCI is successfully leveraging its network management system (NMS) capabilities to increase its overall share in the long-distance market. If MCI can maintain this momentum, we believe the company can increase its growth rate substantially, and thereby increase its share of the long-distance market from 10 percent in 1988 to as much as 16 percent in 1993. Perhaps most important for MCI, success will come from service capabilities, not price.

MARKET POSITION

MCI trails only AT&T in the long-distance market, having captured 10 percent of total revenues in 1988, up from 8 percent in 1987. A 28 percent jump in revenues to \$5 billion can be accounted for largely by MCI's penetration into the *Fortune* 1000 market.

During the past year, MCI has convinced its large accounts that it can meet their needs. Credit for the renewed confidence in MCI goes primarily to its position as the first interexchange carrier to market a sophisticated network management interface.

MCI approaches network management from several vantage points. These are described in the following paragraphs.

Integrated Network Management Services (INMS): INMS is MCI's umbrella term for the network management features it will introduce for its customers.

INMS also identifies its first product, an intelligent, 386-based workstation that provides direct access via a standalone interface to the MCI network on a dial-up 9.6K bps line. The INMS workstation displays network information graphically via proprietary software. INMS permits network monitoring, trouble management, configuration management, reconfiguration, report generation, billing report generation, and order entry/tracking for Vnet and 800, Prism, and TDS 1.5 services.

MCI View: Another INMS application, this service allows MCI customers to use IBM's NetView to monitor their networks. It was introduced commercially in April 1989; test marketing began in mid-1988. MCI View currently works with Vnet and TDS 1.5 services.

In addition to NetView, MCI customers must have 9.6K bps dedicated private lines, a modem, a front-end processor, and program software.

AT&T's Accumaster Integrator, a much more powerful network manager, will become available in the fourth quarter, 1989. Compatibility between Accumaster Integrator and Cincom's Net/Master (a NetView substitute) will also be available then. A simpler, standalone management system for T1 lines, ACCUNET T1.5 Information Manager, was introduced in April 1989.

MCI Network Management Strategy

Sprint will introduce its first proprietary NMS in the third quarter but will not achieve NetView compatibility until 1990.

Vnet: Several basic network management features are embedded in Vnet, MCI's virtual private network service analogous to AT&T's Software Defined Network. These features include Network Information Management System (NIMS) traffic analysis reports and Customer Interface Management System (CIMS), which allows the customer to directly control numbering, feature access, and user code plans. NIMS, introduced last year, and CIMS, introduced two years ago, also fall under the INMS umbrella.

MCI Fax: A dedicated facsimile transmission network, MCI Fax permits greater control over fax communications. Features include message store-and-forward, user codes, and detailed reporting. MCI plans to add inbound message storage, but no release date has been announced. MCI Fax became commercially available in November 1988.

Network Management Revenues

Figure 1 shows MCI's network management estimated revenues. Total revenues are projected to increase from \$5.0 billion in 1988 to \$12.5 billion in 1993, a compound annual growth rate of 20 percent. During


this period, incremental growth from network management activities is projected to rise from \$495.0 million in 1988 to \$3.4 billion in 1993, an annual growth rate of 47 percent. It should be noted that these network management revenues are not direct billings for NM services but represent the increase in total sales stimulated by the availability of those NM services.

By contrast, over this same period, we expect the total market to grow by 9 percent per year, AT&T by 7 percent per year, and Sprint by 20 percent per year.

Therefore, we make the following assumptions:

- MCI's network management capabilities boost growth between 1988 and 1993 from 15 percent to 20 percent per year;
- MCI continues to offer network management features as loss leaders for its other services; and
- network customers now indicating a strong preference for IXC network management solutions and a high opinion of MCI's capabilities in this area have their expectations met.

Figure 1. MCI Network management revenues (in millions).



Source: NBI/Datapro estimates.

MCI Network Management Strategy

PRODUCT/SERVICE STRATEGY

MCI has a well-defined plan for providing a high level of network management capability to its biggest customers, subscribers to Vnet and other services. For smaller companies that spend a disproportionately high amount of their telecommunications budgets on fax, MCI Fax provides a more targeted solution.

MCI developed the INMS and MCI View products in response to requests from several large individual customers, then turned them into generic products for general distribution.

By making its own network NetView compatible, MCI has shifted development costs from itself to IBM and avoided the proprietary protocols that limit the marketability of NMS.

MCI's competitors are using a different approach. Sprint's first network management product, scheduled for release later this year, will be proprietary; NetView compatibility is planned for 1990.

AT&T has more ambitious and complex plans. Its Accumaster Integrator enables customers to control all of their networks at once, whether they are voice, data, packet, AT&T, local phone, NetView, OSI, or any combination of these. MCI View will manage only a customer's MCI links. The Accumaster Integrator will be more difficult to sell and implement, but its primary purpose will be to generate revenues, not to stimulate demand for AT&T's long-distance services.

The ACCUNET T1.5 Information Manager is more competitive with MCI View but does not have an interface to an end-to-end NMS like NetView.

In March 1989, implementation of SS#7 was completed on the entire MCI network. From a technical point of view, MCI now has the capability to expand network control dramatically for its customers.

SALES STRATEGY

Network management is the key to MCI's marketing strategy for large customers. Our research shows that customers are buying the message—large users, and particularly NetView users, are looking to their IXCs for network management capabilities. The 50 to 75 Vnet customers are the primary target for network management services, but the market potential for this product will expand: price cuts are making Vnet affordable to customers with \$25,000 per month phone bills. In addition, INMS features will be available on all 800, Prism, and T1 features by the end of the year.

Nevertheless, with a relatively small number of potential customers and relatively low pricing, the INMS services will generate significant revenues themselves.

The INMS features of Vnet are nominally chargeable, but currently these charges are waived if the customer makes a one-year commitment. The charges are shown in Table 1.

MCI View and INMS terminal prices have not been released, but customers will incur an installation charge plus a monthly fee. Prices are expected to fall within the range of CIMS and NIMS charges.

MCI Fax includes network control features at standard long-distance rates. Although the rates are essentially the same, the rate structure, in terms of time periods and mileage bands, is simpler than that of MCI's standard "Dial 1" service. Calls are also charged for the first 30 seconds (roughly half a page) plus each additional 6 seconds; Dial 1 calls are charged for whole minutes.

Network management will most likely remain a way for MCI to boost its market share for long-distance services, rather than become a profit center in itself. Since the revenue MCI generates from NMS charges are minimal or nonexistent, the company must make up its costs in increased revenues.

STRENGTHS

Alliance with NetView: NetView is the de facto standard in network management (and promises to remain so, at least for the next two or three years). This alliance simplifies network management for MCI and its customers. NetView's focus on data networks (rather than voice) is a major weakness of IBM's product, and MCI has solved this problem—creating a new market for itself.

Solid image, compared to IBM, Digital, and the RHCs: In our recent survey of MIS and telecom managers, we asked them to indicate their overall evaluation of several competing companies as providers of network management systems. The percentage of respondents

Feature	Installation Charge (\$)	Monthly Fee (\$)
CIMS	\$500	\$100
NIMS	950	950
SMDR codes	400	400

Table 1. MCI charges for network management capabilities.

MCI Network Management Strategy

Company	Percentage
AT&T	77
Digital	49
IBM	59
LEC	40
MCI	44
Sprint	29

Table 2. NBI survey results—percentage of respondents rating each company as “one of the best” or “above average”.

rating these companies “one of the best” or “above average” is presented in Table 2.

Strong year among large accounts: MCI had a bang-up 1988, growing more than 25 percent to over \$5 billion in revenues. Growth was strongest among large accounts. MCI's gain among large accounts (as well as Sprint's) came at the expense of AT&T. Our survey of MIS and telecom managers revealed that an above-average number of NetView users (who are committed NMS buyers), as well as organizations that place a high value on single-sourcing, rank MCI as “one of the best” providers of network management.

WEAKNESSES

Mediocre image compared to AT&T: As indicated in Table 2, AT&T is head and shoulders above its competition (including MCI) with respect to the customer's evaluation of its network management systems capability. MCI must continue to grow in order to close this gap.

Ambitious plans: Although MCI has gotten off to a great start in network management compared to the other major carriers, keeping this momentum will not be easy, particularly since MCI's business was built on low cost, not richness of features.

Tepid response to INMS: Despite aggressive pricing, MCI's customers are not clamoring for its INMS fea-

tures. MCI has some work to do to get the formula right.

OSI migration plans: MCI has not committed itself to OSI compatibility for its NMS products, although the company claims to be looking at OSI alternatives. Aligning itself with NetView has its advantages, but IBM's plans for making NetView OSI compatible are sketchy at best. Many of MCI's customers need to know what the migration to OSI will involve.

FUTURE DIRECTIONS

During the past 18 months, MCI has taken share from AT&T at a brisk pace. The large accounts captured by MCI are looking for better network management to reduce costs and to improve their competitive advantage.

MCI has shown leadership in network management, and its sales success demonstrates that its customers believe that this momentum can be maintained.

To bring in new national accounts, MCI must convince its customers that it can reduce their costs and increase their competitive advantage through better network control. While direct revenues from NMS will be limited, MCI's NMS features will be the key to the company's continuing success.

MCI's strategy is completely different from AT&T's. While AT&T is developing the network management market itself with its powerful Accumaster Integrator, MCI is using network management as a loss leader for its long-distance services.

We believe MCI is taking the right course. Our research shows that nearly three quarters of NetView users and more than one half of large users (over \$100 million in revenues) plan to use virtual private networks. By making its own virtual private network (VPN) service NetView compatible, MCI has clearly responded to the market. □

||||| Research Report

Network Equipment Technologies Network Management Strategy

This report will help you to:

- Examine the relationship between N.E.T.'s network management offerings and its T1 network business.
- Evaluate N.E.T.'s network management systems.
- Anticipate N.E.T.'s network management strategies as the demand for T1 products slows over the next few years.

Network Equipment Technologies (N.E.T.) is investing heavily in Network Management Systems (NMS) to generate T1 equipment sales. This strategy may become a serious problem for N.E.T. as the T1 market matures.

MARKET POSITIONS

N.E.T. is the leading supplier of private T1 networks, with approximately 1,500 T1 nodes in service at the end of 1988 at over 100 customer sites.

N.E.T. is a major player in NMS by virtue of its T1 market position. The company sells the following NMS products for its IDNX T1 multiplex systems and other, lower speed products:

Series 5000 Network Management System: a UNIX-based product that runs on a Sun Microsystems terminal, employing color graphics and a relational database for real-time management of N.E.T. T1 networks. Additional workstations can be attached to the Series 5000 via a LAN to provide multiple concurrent operations and applications. This will be N.E.T.'s top-of-the-line system, and will be available in this year's third quarter. The product is priced between \$11,000 and \$50,000, depending on the network configuration.

Series 4010 NMS: a centralized, graphically oriented, UNIX-driven system that runs on a Sun Microsystems

workstation. Introduced in 1988, this is currently N.E.T.'s most powerful NMS for large T1 networks.

Series 3040 Enhanced Operator Console: a proprietary system running on a PC AT or PS/2 in MS-DOS, for basic management of smaller T1 networks. Another version, called the Series 3050 EOC Plus, is available for networks with sub-T1 rate multiplexers.

Series 3210 Network System Interface: allows centralized, host-based management of N.E.T. T1 networks with IBM's NetView/PC.

IBM Transmission Network Manager: a software product that runs on a PS/2 Model 80 in OS/2. This system brings SNA NetView control to IDNX systems resold by IBM.




Figure 1. Forecast of N.E.T. network management revenues (\$ thousands).

Network Equipment Technologies Network Management Strategy

In addition to these products, in April 1989 N.E.T. introduced two network management services to T1 product customers with maintenance contracts:

Expert T-span Service: N.E.T. takes end-to-end responsibility for network up-time, coordinating the solution of N.E.T. equipment and carrier service problems for the customer. N.E.T. service contract customers pay an additional \$180 per month per T-span for nonintelligent multiplexers and \$148 for intelligent ones.

Expert Fault Management Service: a real-time, artificial intelligence-based system that identifies network faults, diagnoses the problem, and forwards trouble tickets to N.E.T. for resolution. The product runs on a Sun Microsystems workstation, and works in conjunction with the Series 5000 NMS. The product's first installation is expected to take place in June 1989, and the service charge for a typical small network of three-to-five nodes would be approximately \$25,000 per year.

These services, and maintenance in general, are provided by the N.E.T. Technical Assistance Center (TAC), with some 40 engineers and support staff; technicians are dispatched by the TAC from regional locations.

Network Management Revenues

N.E.T.'s direct revenues from NMS will be limited. Indeed, N.E.T. does not even expect to recoup its research and development costs from direct revenues. Rather, the company views NMS as a marketing expense, to sell more T1 multiplexers.

Figure 1 shows direct NMS revenues, based on the penetration of existing accounts. By 1993, we project that 50 percent of N.E.T.'s customers will have NMS installed, generating a mere \$4.1 million in direct revenues.

Looking at NMS another way, however, a very high proportion of N.E.T.'s total revenues over the next five years could be attributed to its NMS activities. N.E.T. is differentiating its products in two ways: better network management capabilities and better customer service.

Our observations are based on the following assumptions:

- growth of T1 systems installed begins to level off after 1990;

- penetration of the installed T1 base with NMS grows rapidly;
- as smaller T1 networks are supplied with NMS, the average cost per system declines; and
- growth in NMS sales slows precipitously after 1990.

PRODUCT/SERVICE STRATEGY

To sell more T1 systems, N.E.T. is developing NMS to operate in both the OSI and SNA environments. In so doing, N.E.T. has its bases covered, no matter which way the market turns.

By making its products either OSI-compatible or SNA-compatible, N.E.T. opens up opportunities in *multivendor* environments. N.E.T.'s long term strategy is to capture an increasing share of its customers' expenditures on their networks; thus, N.E.T. products must fit easily into multivendor applications.

Although the Series 5000, 4010, 3050, and 3040 products are now proprietary, N.E.T. plans to make them OSI-compatible, as those standards develop. This approach ensures that they work with products from other vendors, and also simplifies development decisions.

Similarly, the use of NetView standards will allow N.E.T. to shift the burden of product development to IBM, also N.E.T.'s distributor.

With the introduction of network management services, N.E.T. is widening the net to capture customer budgets even further. NBI research shows that users, especially those with NetView, are looking seriously at facilities management alternatives; if the shift away from "do-it-yourself" networks has begun, N.E.T.'s introduction of its Expert T-span Service is most timely.

As the do-it-yourself era peaks, users are also looking to their carriers for network management solutions. To explore public network opportunities, in April 1989 N.E.T. announced a joint venture with Tellabs, a strong player in the carrier market. Such a combination will help protect N.E.T.'s flank should its customers start looking to the RHCs and IXC's for virtual T1 networks.

SALES STRATEGY

N.E.T.'s sales strategy is strictly national accounts, with a three-level approach:

Network Equipment Technologies Network Management Strategy

Market	Percent
Manufacturing	31
Financial	21
Utilities/communications	12
Banking	11
Government/defense	7
Retail	6
Transportation	6
Other	6

Table 1. N.E.T. sales by vertical markets for fiscal year 1988.

- **direct sales:** with its own sales force of some 85 salesmen, N.E.T. is targeting the largest 600 private networks in the USA;
- **IBM sales:** through its OEM deal with IBM, N.E.T. is reaching the SNA market, a large part of which overlaps with its direct sales territory; and
- **international distribution:** outside the USA, N.E.T. will use distributors (including IBM) to reach the largest 400 private networks; international sales to date are limited, but a distribution agreement with British Telecom International was announced in April.

Despite its meteoric growth, half of N.E.T.'s sales are still coming from new accounts. N.E.T. reckons that only 12 percent of the "N.E.T. 1000" (the 1000 largest users of telecom networks in the world) are now customers. Lots of room for growth remains.

N.E.T.'s sales by vertical market for fiscal year 1988 are shown in Table 1.

During the next few years, growth in T1 demand will slow precipitously as the process of converting analog private line users to T1 digital service is completed. At that point, the T1 equipment market will be good, but nothing like it is today.

To take up the slack, N.E.T. is looking in three directions: sales through IBM, international markets, and higher (DS3) and lower (DS0) rate systems.

STRENGTHS

T1 market position: N.E.T. is in a strong position in T1, and is likely to garner a large share of the high-level NMS expenditures of its accounts; in 1988, N.E.T. took the number one spot in T1 away from Timeplex.

IBM partnership: N.E.T.'s relationship with IBM ensures the company an entree into SNA accounts for T1 and NMS systems.

Tellabs partnership: as the market turns to virtual private networks, N.E.T. can try to capture some of the lost revenues by working with Tellabs, a strong public network market player.

OSI migration strategy: while making a strong play for the SNA market with IBM, N.E.T. has wisely committed to OSI compatibility; its bases are covered.

Strong image: for a small company, N.E.T. has a strong image among network management users. While only 12 percent of all respondents to a recent NBI survey ranked N.E.T. as "one of the best" or "above average" as a provider of NMS, 15 percent of larger companies, 15 percent of current NMS users, 15 percent of those currently in a buying mode, and 25 percent of those with NetView ranked N.E.T. as one of the best or above average.

Highly intelligent T1 systems: by embedding a high level of network management capability into its intelligent T1 muxes, N.E.T. has also shut out other suppliers from lower level NMS sales.

WEAKNESSES

T1 market tapped out: during the next few years, growth in T1 demand will slow precipitously as the process of converting analog private line users to T1 digital service winds down. At that point, the T1 equipment market will be good, but nothing like it is today.

Limited direct revenues from NMS: An NMS is now considered a marketing cost for T1 systems; the company does not expect NMS revenues to cover R&D costs. Can a \$137 million company afford to make R&D investments that don't pay off right away? The answer is no. While \$16 million (12 percent of revenues) was expensed for R&D in fiscal year 1989, an additional \$13 million was carried on the books as capitalized software development costs, up \$8 million from the year before. Essentially all NMS development is for software; is N.E.T. capitalizing what is, in effect, marketing expense?

FUTURE DIRECTIONS

N.E.T. has clearly charted its course in network management, and the company is reaping the rewards in T1 sales today. With high-powered NMS for both OSI and SNA applications, N.E.T. has covered the market;

Network Equipment Technologies **Network Management Strategy**

by aligning itself with IBM, it has reduced the cost of NMS development, but the costs remain high. As long as N.E.T. can recoup these costs through T1 hardware

profits, the company will remain the number one up-and-comer in network management. □

Northern Telecom Network Management Strategy

This report will help you to:

- Pinpoint the market segments to which Northern Telecom provides network management systems.
- Evaluate Northern Telecom's network management product/service strategy.
- Judge the company's capability to bring unique network management products to its customers.

Northern Telecom (NTI) seems to have ambitious goals for network management, on both the public and private sides of its business. But success will be elusive—with limited commitment and poor execution, NTI has little to tempt customers.

POSITION

Northern Telecom is the leading supplier of central office (CO) and PBX equipment in the world. But NTI fails to match its success as the low-cost supplier of telecom hardware in the higher value-added end of the market: providing software and service-based solutions for managing the networks of its large business and telephone company customers.

Northern Telecom provides network management systems (NMSs) to four distinct market segments:

Operations systems: NTI has several operations, maintenance, and administration systems for telephone company application, including a centralized automated loop reporting system (introduced in 1976), a variety of remote and local test systems, and its Digital Facility Management System (DFMS), introduced in 1984.

In addition, the company has three carrier network management applications that run on its DNC-500 controller: Transport Services Management, which

provides automated and centralized control and management of telco T1 networks; Dynamically Controlled Routing, which performs alternate routing management functions; and Billing Services Management, which processes automated message accounting data. This family of DNC-500-based systems was introduced in 1986.

Centrex management systems: an additional application for the DNC-500 is Business Network Management, which enables telcos to provide Centrex customers with station moves and changes control and other network management features.

Data network management systems: "Meridian Network Control DNS Network Manager" is a Sun Microsystems terminal programmed in UNIX for




Figure 1. Northern Telecom's network management revenues in millions.

Northern Telecom Network Management Strategy

managing the "Meridian Data Networking System," introduced in October 1988; in essence, NTI is setting itself up to compete against IBM; its network manager serves the function of IBM's NetView.

Systems integrator: in alliance with Hewlett-Packard, NTI established its Corporate Networks Operation (CNO) in 1988 to help large HP and NTI accounts build multivendor networks.

Network Management Revenues

NTI revenues from its NMS are estimated at \$66 million in 1988. We project that they will reach \$105 million by 1993. While some two thirds of revenues are now derived from public network applications, with the balance from private applications, the proportions will shift by the end of the forecast period. Private applications will grow by the sheer force of NTI's PBX market position, accounting for the bulk of NMS sales in 1993. Public network sales will languish at their current levels, which represent about 1 percent of NTI's total revenues.

Our analysis of NTI makes the following assumptions:

- ten percent of data networking systems revenues are for network management systems;
- no break is made on the OS front; and
- figures are for worldwide sales in U.S. dollars.

PRODUCT/SERVICE STRATEGY

Northern's product strategy has been fragmented so far, yielding poor results. Public network products designed for Bell Canada, including the DNC-500 family, have not been well received by the regional holding companies (RHCs); AT&T has a lock on the Operational Support Systems (OSS) market, and NTI has done little to scare its competitor, despite the Canadian company's enormous success in digital switching.

Northern's OSS strategy is to push OSI standards in order to open up AT&T's installed base. Unfortunately for NTI, its OSI-based Network Operating Protocol is not compatible with the AT&T systems in which the RHCs have already invested billions. At the same time, AT&T has also adopted OSI standards for its OSS systems, retaining the high ground in this market. OSS market difficulties have been compounded by in-fighting between incompatible products. The DFMS is built around a Digital Equipment

Corporation VAX, while the DNCs use NTI's DV-1 technology. The company has not decided which approach to use in the future.

On the private network side, its customers are spending big on network management, but fragmented and overlapping distribution channels hinder Northern's capability to maintain the account control necessary to sell network management systems.

NTI's first stab at the market for private users was to introduce the DNC family of products, giving network management capabilities to virtual private network users. This was a great idea in search of a market. The telcos distributing the systems were even worse at account management than NTI.

The second stab at the market attempts to get NTI into IBM shops with a NetView type of product. While the product will have a number of feature improvements over NetView, NTI's strategy is to sell the system through the RHCs. The direct PBX sales force does not want to touch the product. According to one NTI sales manager, "I don't want my guys selling Meridian Data Networking System—it doesn't work."

Without a position in the corporate T1 market, Northern is competitively disadvantaged. Network management opportunities are going to others, such as N.E.T., which dominates the T1 market. Companies are not building network management around PBXs.

Without account control, product development is doomed to failure. The RHCs and other interconnects compete with NTI's direct sales in a confusing and demoralizing fashion. When the SL-1 wins the bid, NTI's distribution channels then start competing with each other on price. Money for such value-added applications as network management remains on the table.

NTI has big plans for systems integration and even facilities management, but until it gets control of distribution and account management, these plans will not materialize. The first step in this direction, the CNO joint venture with HP, is not breaking any records.

While pushing OSI standards on the public front to crack AT&T's lock on the OS market, Northern is emphasizing SNA compatibility on the private side to get a piece of the IBM base. While this is a sound approach, it's failing on execution.

Northern Telecom Network Management Strategy

SALES STRATEGY

OS and Centrex management systems are targeted to central office and transmission accounts.

Data network systems are, for the time being, sold to non-NTI accounts by a small, dedicated sales force recruited largely from IBM. However, within two to three years, NTI says these products will be sold through the RHCs.

PBX sales are made through at least four channels: direct, RHCs, independent telcos (e.g., Centel), and interconnects. For large national accounts (Premier Accounts in NTI nomenclature), a special team is supposed to coordinate all sales activities. In reality, overlapping and competing sales forces spell lack of account control. Northern competes against itself on price. Value-added applications are supplied by others.

The Corporate Networks Operation (CNO) will sell to existing HP and NTI accounts. CNO will charge for its services and says it will recommend the best equipment for the job; however, the organization seeks to thereby sell its own existing products.

STRENGTHS

Strong base of accounts: among carriers and large private accounts, NTI has an exceptionally strong client list.

Public network experience: NTI is in a strong technical position for developing innovative operations systems; with Bell Canada, NTI has a test bed for these products—why can't the company deliver?

WEAKNESSES

Locked out of the operational support systems market: in the telco market, AT&T sets the standards,

and opportunities for selling into the AT&T base are few; with a limited product line and competing products, NTI's future looks grim in this market.

Poor differentiation: On the data network side, NTI's new products are poorly differentiated, considering that the company is a newcomer to the market. Despite some technical improvements over NetView, there appears to be no compelling reason for customers to turn to NTI. Perhaps this problem can be shifted to the RHCs.

Products are technology driven: Not surprisingly, NTI's private voice network systems have been poor performers because they were designed for Bell Canada and its telco requirements; the needs of the Centrex customer have not been met.

No T1 market position: Network management focuses on the transmission links, not the switches. NTI does not compete in the private T1 market.

Poor sales force morale: Forced to compete with other channels on price, NTI's sales force is demoralized. The problem is aggravated by the lack of products for corporate networks.

FUTURE DIRECTIONS

NTI poses little threat to its competitors in network management, either in the public or private network markets. For its customers, the company offers little other than switches.

Given the enormous resources the company must invest in central office and PBX enhancements, no major breakthroughs in NMS are anticipated from NTI. Should it develop a product, problems with distribution might prevent reaching the right customers. □

NYNEX Network Management Strategy

This report will help you to:

- Assess the impact of NYNEX' acquisitions on its network management strategy.
- Evaluate the effectiveness of NYNEX' OSS investments in making network control available to its telephone company customers.
- Compare NYNEX' success with the track record of other Bell Operating Companies.

MARKET POSITION

NYNEX' first steps to bring better network control to telco customers have been halting. Through a deregulated subsidiary, however, NYNEX has built—largely by acquisition—a \$300 million network management business. Internal growth will prove more difficult.

NYNEX' operating companies, New York Telephone and New England Telephone, offer two types of network management features to Centrex customers:

1AESS Centrex: CLAS (over 200 lines) and Mini-CLAS (under 200 lines) allow the customer to do moves and changes either from a dumb terminal or, in the case of MiniCLAS, with a touch-tone phone. Changes, which are not done in realtime but the next business day, can include feature assignment and line restrictions. Reports are available for CLAS.


Digital Centrex: This service features automatic route selection, station message detail recording, authorization codes, and customer service administration (including station and feature changes).

V-Path Customer Network Service: This service provides a virtual private wide area network for intra-LATA use. Features include uniform numbering, authorization codes, and customer-controlled moves and changes; and are available from all equal access-conforming SPC central offices.

Intellihub Dedicated Network Service: This is a "pre-ISDN" virtual private intraLATA WAN with two key features beyond those available with V-Path: T1 links within the networks and customer control of traffic routing and network reconfiguration. Intellihub is available only from central offices equipped with Northern Telecom DMS-100 switches.

Through its unregulated subsidiary NYNEX Information Solutions Group, Inc. (ISG), NYNEX has five operating companies active in network management:

1. *AGS Computers, Inc.:* This company provides project management services and software and data



Source: Northern Business Information/Datapro.
Figure 1. NYNEX network management revenues—unregulated activities (in millions).

NYNEX Network Management Strategy

communications software, particularly for financial concerns. NYNEX acquired AGS Computers in October 1988. AGS currently has 2,800 employees.

2. *Telco Research Corp.*: This company provides call accounting and traffic management systems. Acquired in April 1986, the company has 130 employees.

3. *NYNEX Computer Services*: This group provides systems integration and facilities management services, primarily to other NYNEX companies; it has 200 employees.

4. *NYNEX Development Company*: This operating company nurtures strategic new business opportunities, including network management control projects; it was formed in January 1985.

5. *Complex Systems Integration Group*: This organization provides systems integration for large telecom and computer projects; it was formed in January 1989.

Note: Only network management-related activities are described above.

Estimates of network management revenues are shown in Figure 1. The revenue estimates do not include any future acquisitions and are based on the following assumptions:


- NYNEX is able to absorb AGS successfully
- Telco Research products are extended beyond call accounting to data network management systems (NMSs)
- AGS grows at 8 percent annually; Telco Research grows 10 percent annually; Computer Services grows 15 percent annually

PRODUCT/SERVICE STRATEGY

NYNEX' corporate strategy includes three goals: build its core telco business, expand into the computer market, and expand internationally. Network management will help NYNEX with the first two of these three objectives.

NYNEX has cast a wide net beyond its telephone operations, using ISG as the vehicle for expansion into computers. NBI/Datapro estimates that over half these activities may be considered network management related, as described above.

ISG doubled in size last year with the acquisition of AGS. Growing these acquired businesses and giving them clear focus will not be easy. AGS itself has seven



Source: Northern Business Information/Datapro.

Figure 2. NYNEX OSS investment.

operating units. Nevertheless, NYNEX now has all the bases covered in network management, including NMS products, facilities management, and systems integration.

On the regulated side, NYNEX' approach is driven largely by what is available from its suppliers but constrained by regulation. Advanced features for Centrex and virtual private networks offered by Northern Telecom and AT&T for their electronic switches have been tariffed, but regulation limits NYNEX' ability to offer interLATA service and competitive pricing. Until recently, regulation also limited a Regional Holding Company's ability to offer "one-stop shopping" for equipment and service. That constraint has been removed.

NYNEX' objective is to make network control available to its telephone company customers. To achieve this goal, the company is investing heavily: NYNEX is currently installing digital lines faster than any other carrier in the U.S. By 1993, 100 percent of its network will be electronic; 74 percent will be digital.

To manage its own network, NYNEX is combining cost reduction efforts with new investment in operations support systems (OSS). Figure 2 depicts that investment, which includes in OSS: FACS, COSMOS, DFMS, SARTS, and TIRKS systems.

SALES STRATEGY

NYNEX does not have a unified profile for its network management capabilities.

For large-scale systems integration opportunities, the Complex Systems Integration Group (part of ISG) will act as the account coordinator for unregulated activities. On the regulated side, NYNEX Systems Marketing provides one-stop shopping for large accounts. In

NYNEX Network Management Strategy

short, it is difficult, even for its biggest accounts, to take advantage of all of NYNEX' capabilities without multiple points of contact.

This complexity is illustrated by a contract that NYNEX won in 1988. NYNEX is lead vendor in an 8,000-line contract from New York State (Empire Net). Under this contract:

- N.Y. Tel provides service
- NYNEX Business Information Systems provides CPE and maintenance
- ISG provides project management, systems integration, network management, and software through the Complex Systems Integration Group, NYNEX Computer Services, and Telco Research
- NYNEX Systems Marketing acted as sales agent

Each of ISG's units maintains its own sales force. Telco services are sold by the individual operating companies; NYNEX Business Information Systems Company is also an agent for Centrex.

Loss of the General Electric account to AT&T Tariff 12 indicates the problem, caused by regulation, which prevents NYNEX from responding flexibly to large customers with major network management requirements. But NYNEX is learning. In January, for example, N.Y. Tel won a major contract from the Securities Industry Association in New York, beating out a number of competitors, including AT&T and Teleport.

STRENGTHS

Concentration of Big Business: The sheer number of corporate headquarters located in New York gives

NYNEX an advantage in selling network management, decisions about which tend to be centralized.

ISDN Strategy: Aggressive pricing of the second "B" channel for basic ISDN service will go a long way toward ensuring customer acceptance. Widespread use of ISDN will be the best way for NYNEX to improve network management for its customers.

Management of Large Accounts: The establishment of the Complex Systems Integration Group will help coordinate activities for large accounts, but its effectiveness without carrier services in its portfolio may be greatly reduced (see "Weaknesses" below).

WEAKNESSES

Fragmented Profile: NYNEX' profile to the customer is fragmented; NYNEX Systems Marketing, NYNEX Business Information Systems Company, and NYNEX Information Solutions Group target many of the same customers and all of New York and New England Tel's biggest accounts. "Integrated" solutions cannot be provided by three organizations.

Marketing Control: Problems of overselling customer calling features (i.e., stimulating demand in areas where services are not available) could indicate difficulty in controlling marketing; custom calling features are, after all, a form of network management.

Regulatory Constraints: MFJ restraints limit NYNEX' ability to solve its "fragmented profile" problem.

Given the limits imposed by the MFJ, NYNEX is probably doing the best it can in network management: acquiring expertise, but keeping these organizations distinct. However, until its large accounts group, NYNEX Systems Marketing, can bring all capabilities to bear for the customer *and* price competitively, NYNEX' real capabilities will be untested. □

US Sprint Network Management Strategy

This report will help you to:

- Examine what US Sprint is doing to bring network management capabilities to its customers.
- Discover how US Sprint's INSITE network management offering fits into its overall strategy.
- Determine whether Sprint's network management offerings can benefit your network operations, now and in the future.

US Sprint's plans for making network management information and control available to its customers are lagging behind those of its key competitors, AT&T and MCI. Sprint will have to pick up the pace to stay competitive in this vital area.

With what is possibly the most advanced network in the business—it is 100 percent digital and has SS#7 fully implemented—Sprint is moving to bring some of its network management capabilities to its customers. These efforts include:

MARKET POSITION

US Sprint, a partnership between United Telecommunications, Inc. and GTE Corp., is the number three long-distance carrier. It garnered 6.7 percent of the long-distance market in 1988, up from 5.3 percent in 1987. Revenues grew by 36 percent in 1988 to over \$3.4 billion. In the first quarter of 1989, revenues grew by 30 percent, and the company was solidly in the black. Clearly, Sprint's outlook is bright.

Sprint's subsidiary, Telenet Communications Corp., is the leading value-added network provider in the United States, with a 45 percent market share. Sales were up by 20 percent in 1988 to \$216 million.

In 1985, US Sprint became the first carrier to introduce a virtual private network (VPN), and today the company has approximately 150 to 200 VPN customers. Sprint claims to be the primary carrier for 110 out of the top 800 telecommunications users in the United States.




Figure 1. US Sprint network management revenues in millions.

US Sprint Network Management Strategy

INSITE (Integrated Network System Interface and Terminal Equipment): This proprietary software package permits customer monitoring and trouble ticketing.

Virtual Private Network: Several basic network management features are embedded in VPN, Sprint's virtual private network service that is analogous to AT&T's Software Defined Network. These features include INSITE and customized billing reports.

Management reports and call detail on tape: All Sprint's services, including MTS, WATS, and in-WATS include detailed billing reports and optional call detail on magnetic tape.

Telenet: In addition to operating its public VAN network, Telenet designs, implements, and operates private and virtual private packet networks. These private data networks can be built with essentially any level of network management that the customer requires. The TPS5 Network Management System family ranges from the TPS5/II-5001, a microcomputer-based NMS for small or regional networks, to the TPS5/II-5975, a minicomputer-based NMS for networks with up to 200 nodes. A color network monitor offers full graphic displays.

Network Management Revenues

Sprint's NMS estimated revenues are shown in Figure 1. Sprint's total revenues are projected to rise from \$3.4 billion in 1988 to \$8.5 billion in 1993, an annual growth rate of 20 percent. During this period, incremental growth accounted for by network management activities is projected to rise from \$149 million in 1989 to \$1.7 billion in 1993. We believe these activities will increase total revenue growth by 2 percent per year. Note: these NMS revenues are not direct billings for NM services, but rather the increase in total sales stimulated by the availability of those NM services.

By contrast, over this period, we expect the total long-distance market to grow by 9 percent per year, AT&T by 7 percent per year, and MCI by 20 percent per year.

PRODUCT/SERVICE STRATEGY

INSITE, primarily for VPN users, is the cornerstone of Sprint's network management plans. The system consists of a PC loaded with terminal emulation software, connected via Telenet to a Sprint host at its

Network Operations and Control Center (NOCC) in Kansas City. Application and customer information resides at the Sprint host.

From any point on its network, a VPN customer can access traffic data that is updated hourly; enter and track trouble tickets; and activate and deactivate credit cards and off-net codes.

Information is presented only as text. No graphics are available.

The initial version of the system is now in beta test, with commercial availability planned for third-quarter 1989. There are three beta test sites.

Version Two, although not now scheduled, will probably be introduced in 1990. This version will include several major improvements: graphics display capability; NetView-compatible interface; and more data and control from Sprint's own NOCC system.

Sprint is already experimenting with more sophisticated network management for the government under the FTS-2000 contract.

With Version Two, most of the application software will reside in the terminal; therefore, a PC with more horsepower, possibly a PS/2, will be required. Sprint is less inclined to use a Sun workstation, but a final decision has not been made.

Sprint's own network is managed from the NOCC and two Regional Control Centers in Sacramento and Atlanta, connected to its 43 digital switches (41 DMS-250s and 2 DMS-300s) using 70 SS#7 nodes.

SALES STRATEGY

Sprint now focuses its network management efforts on its largest VPN customers. INSITE will target Sprint's very largest users, and the company expects to get 50 customers for Version One. Pricing will be sufficiently attractive for Sprint to believe that most of its target group will sign up. INSITE is not a profit or revenue center; rather it is priced to enhance VPN sales.

Final prices have not been set, but pricing for the first version will be based on an installation plus a monthly fixed charge. Version Two will probably add some kind of usage and event charges.

Telenet targets the same *Fortune* 500 market as Sprint but also has a stronger presence internationally. The company has developed and installed over 100 private data networks and more than 20 elec-

US Sprint Network Management Strategy

tronic mail systems worldwide. The company also claims to have more than 2,000 virtual private networks supported on its public data network.

In December, Telenet moved down-market with the introduction of an entry-level PAD (packet assembler/disassembler) and dedicated access facility package. In the first five months, about 24 companies purchased these entry-level hybrid systems. Telenet wants to become a major supplier of virtual private networks to small- and medium-sized businesses.

Telenet and Sprint maintain separate sales forces and operate essentially independently. The overlap of their account lists, built up when they were separate, is limited. More important, buyers of voice and buyers of data within a customer account remain distinct, according to the company. A systems integration group at Telenet, however, has been charged with coordinating account management between Sprint and Telenet. Its first accounts are FTS-2000 and the State of Illinois.

STRENGTHS

State-of-the-art network: Sprint has a modern network with all-fiber transmission and all-digital switching. The network uses all Northern Telecom switches and has SS#7 fully implemented. Technologically, Sprint is probably the carrier in the best position to open up network management to its customers.

Strong year among large accounts: Sprint grew by more than 30 percent in 1988 to \$3.4 billion in revenues; growth was highest among *Fortune* 500 accounts.

FTS-2000 contract: Winning this major contract with the federal government ensures Sprint's long-term viability. In addition, network management features developed for the government will find applications in the business market.

Telenet's VPN and international expertise: Telenet's position as a major player in the world market for private and hybrid public/private data networks gives Sprint a significant advantage, although the company is not yet fully exploiting the opportunity.

WEAKNESSES

Lesser image compared to other IXC's, IBM, Digital Equipment, and the RHCs: In a NBI recent survey of MIS and telecommunications managers, asked, "Please indicate your overall evaluation of each com-

pany ... as providers of network management systems." The percentage of respondents rating companies "one of the best" or "above average" were as follows:

- AT&T—77.3
- Digital Equipment—49.0
- IBM—58.8
- LEC—40.4
- MCI—43.6
- Sprint—29.3

Slow Start in NMS: Compared to the network management capabilities now offered by MCI and AT&T, Sprint has some catching up to do.

No plans for OSI: Sprint has not charted a course for its NM services to OSI compatibility, nor has it committed to do so.

FUTURE DIRECTIONS

Given the advanced state of its own network and the management capabilities of that network, the degree of control extended to customers is surprisingly limited.

INSITE is certain to be a success with Sprint's customers, but these customers want more. Many expect their carriers to be the primary source of network management solutions. Sprint must pick up the pace if its image is to be improved and if network management is to sell VPN.

Although AT&T and MCI are taking different approaches from each other, both vendors have aggressive plans under way for network management. Unless it can differentiate VPN on features, Sprint will not become the carrier of choice for large companies and will be forced to sell on price alone.

In summary, we assume that:

- Sprint will improve its image.
- Network customers now indicating a strong preference for IXC network management solutions will overcome their reservations about Sprint.
- Sprint's network management capabilities will increase the revenue growth rate by 10 percent from what it would be otherwise. □

Directory of Suppliers

ADC Telecommunications

4900 W. 78th Street
Minneapolis, MN 55435
(612) 893-3070

Product Types:
Network Management Systems
Test & Monitoring Equipment

Allen-Bradley, Communication Div.

555 Briarwood Circle
Ann Arbor, MI 48104
(313) 668-2500

Product Types:
Local Area Networks

Allnet

30300 Telegraph Road, Suite 350
Birmingham, MI 48010
(313) 647-4060

Product Types:
Interstate Switched Facilities
Private Line Facilities

Altos Computer Systems

2641 Orchard Parkway
San Jose, CA 95134
(408) 946-6700

Product Types:
Local Area Networks

American International Communications Corp.

4760 Walnut St.
Suite 105
Boulder, CO 80301
(303) 444-6675; (800) 821-0528; telex 499
0706 XZZZY

Product Types:
Automatic Routing Software

American Network

915 Main Street
Vancouver, WA 98668-3535
(206) 695-3838

Product Types:
Interstate Switched Facilities
Private Line Facilities

American Private Line Services Inc.

39 Border Street
West Newton, MA 02165
(617) 965-5600

Product Types:
Private Line Facilities

American Telecorp

10 Twin Dolphin Drive
Redwood City, CA 94065
(415) 595-7000

Product Types:
Telemanagement Systems

Ameritech Corp.

30 S. Wacker Drive
Chicago, IL 60606
(312) 750-5000

Product Types:
Interstate Switched Facilities
Private Line Facilities

Amnet, Inc.

1881 Worcester Road
Framingham, MA 01701
(617) 879-6306

Product Types:
Packet Switches

Apollo Computer Inc.

330 Billerica Road
Chelmsford, MA 01824
(617) 256-6600

Product Types:
Local Area Networks

Apple Computer, Inc.

20525 Mariani Avenue
Cupertino, CA 95014
(408) 996-1010

Product Types:
Local Area Networks

Applitek Corp.

107 Audobon Road
Wakefield, MA 01880
(617) 246-4500

Product Types:
Local Area Networks

Asher Technologies, Inc.

1 Quad Way
Norcross, GA 30093
(404) 923-6666

Product Types:
Local Area Networks

AST Research, Inc.

2121 Alton Avenue
Irvine, CA 92714
(714) 863-1333

AST Canada
6549 Mississauga Rd
Mississauga, ON L5N 1A6
(416) 826-7514

Product Types:
Local Area Networks
Micro-to-Host Communications Products

Astrocom Corp.

120 W. Plato Boulevard
St. Paul, MN 55107
(612) 227-8651

Product Types:
Local Area Networks

AT&T

295 N. Maple Avenue
Basking Ridge, NJ 07920
(201) 221-2000

Product Types:
Integrated Voice Data PBXs
Interstate Switched Facilities
Local Area Networks
Network Management Systems
Packet Switches
Private Line Facilities

ATT/Paradyne Corp.

8550 Ulmerton Road
Largo, FL 33540
(813) 530-2000

Paradyne
100 York Boulevard, Suite 200
Richmond Hill, ON L4B 1J3
(416) 494-0453

Product Types:
Network Management Systems
Packet Switches

Atlantic Research Corp.

7401 Boston Boulevard
Springfield, VA 22153
(703) 644-9000

Atlantic Research Corp., Louis Albert As-
soc.

P.O. Box 7160, Vanier Branch
Ottawa, ON K1L 8E3
(613) 748-8918

Product Types:
Network Management Systems
Test & Monitoring Equipment

Directory of Suppliers

Avant-Garde Computing, Inc.

Avant-Garde Computing, Inc.

8000 Commerce Parkway
Mount Laurel, NJ 08054
(609) 778-7000

Product Types:
Network Access Control Devices
Network Management Systems

Avanti Communications Corp.

Aquidneck Industrial Park
Newport, RI 02840
(401) 849-4660

Product Types:
Configuration Management Systems
Multiplexers
Performance Management Systems

Banyan Systems Inc.

115 Flanders Road
Westboro, MA 01581
(617) 898-1000

Product Types:
Local Area Networks

BBN Communications Corp.

70 Fawcett Street
Cambridge, MA 02238
(617) 497-2800

Product Types:
Packet Switches

Bell Atlantic Corp.

1600 Market Street
Philadelphia, PA 19103
(215) 963-6000

Product Types:
Interstate Switched Facilities
Private Line Facilities

Bellcore Licensing

290 West Mt. Pleasant Avenue
Livingston, NJ 07039
(800) 521-CORE (2673)

Product Types:
Network Management Software
Network Management Documentation
Network Management Training Services

BellSouth Corp.

1155 Peachtree Street NE
Atlanta, GA 30367-6000
(404) 292-4000

Product Types:
Interstate Switched Facilities
Private Line Facilities

BGS Systems, Inc.

128 Technology Center
Waltham, MA 02254
(617) 891-0090

Product Types:
Network Management Systems

Brightwork Development

P.O. Box 8728
Red Bank, NJ 07701
(800) 552-9876

Product Types:
LAN Management Software

British Telecom, Ltd.

British Telcom Centre
81 Newgate Street
London EC1A7AJ

Product Types:
International Switched Facilities

Bull HN Information Systems, Inc.

200 Smith St.
Waltham, MA 02154
(617) 895-6000

Product Types:
DNS Network Management Software

Bytex Corp.

Southborough Office Park, 120 Turnpike Road
Southborough, MA 01772-1886
(617) 480-0840

Product Types:
Network Management Systems, electronic matrix switches

Cable & Wireless North America, Inc.

8321 Lemmon Avenue
Dallas, TX 75209
(214) 353-0396

Product Types:
Interstate Switched Facilities
Private Line Facilities

CACI, Inc.

3344 N. Torrey Pines Court
La Jolla, CA 92037
(619) 457-9681

Product Types:
Computer Network Design and Analysis Software
Telecom Network Performance Analysis Software

CNPC Telecommunications

Candle Corp.

1999 S. Bundy Drive
Los Angeles, CA 90025
(213) 207-1400

Product Types:
VTAM Network Management Software

CASE/Datatel

7200 Riverwood Drive
Columbia, MD 21046
(301) 290-7710

Product Types:
Network Management Systems

CASE/Datatel

55 Carnegie Drive
Cherry Hill, NJ 08003
(609) 424-4451

Product Types:
Network Management Systems

Century Analysis, Inc.

114 Center Ave.
Pacheco, CA 94553
(415) 680-7800; telex 910 481 3011

Product Types:
Unix-based Network Management Software

Cincinnati Bell Information Systems, Inc. (CBIS)

P.O. Box 1638
600 Vine St.
Cincinnati, OH 45201
(513) 784-5959; (800) 327-3900

Product Types:
Cable Management Software

Cincom Systems Inc.

2300 Montana Avenue
Cincinnati, OH 45211
(513) 662-2300

Cincom Systems of Canada
130 Dundas Street E., #201
Mississauga, ON L5A 3V8
(416) 279-4220

Product Types:
SNA Network Management Software

CNPC Telecommunications

3300 Bloor Street West
Toronto, ON M8X 2W9
(416) 232-6760

Product Types:
Interstate Switched Facilities
Private Line Facilities

Directory of Suppliers

Codenoll Technology Corp.

Codenoll Technology Corp.

1086 N. Broadway
Yonkers, NY 10701
(914) 965-6300

Product Types:
Local Area Networks

Codex Corp.

Maresfield Farm, 7 Blue Hill River Road
Canton, MA 02021-1092
(617) 364-2000

Motorola Information Systems Ltd.
9445 Airport Road
Brampton, ON L6S 4J3
(416) 793-5700

Product Types:
Local Area Networks
Network Management Systems

Communication Devices Inc.

One Forstmann Court
Clifton, NJ 07011
(201) 772-6997

Product Types:
Modem-based Network Monitoring Software

Communication Machinery Corp.

1421 State Street
Santa Barbara, CA 93101
(805) 968-4262

Product Types:
Local Area Networks

Communications Management Systems

1500 Planning Research Drive
McLean, VA 22102-5095
(703) 556-2300

Product Types:

Communications Solutions Inc.

992 S. Saratoga-Sunnyvale Road
San Jose, CA 95129
(408) 725-1568

Product Types:
SNA Network Management Software

Compco

5120 Paddock Village Ct.
Brentwood, TN 37027
(615) 373-3636

Product Types:
Telemangement Systems

Computer Associates International Inc.

711 Stewart Avenue
Garden City, NY 11530-4787
(516) 227-3300

Product Types:
MVS/VTAM Systems Management Software

Computer Network Technology Corp.

9440 Science Center Drive
New Hope, MN 55428
(612) 535-8111

Product Types:
Local Area Networks

Comsat International Communications Inc.

950 L'Enfant Plaza SW
Washington, DC 20024
(800) 392-9992

Product Types:
Interstate Switched Facilities
Private Line Facilities

Concord Communications Inc.

753 Forest Street
Marlboro, MA 01752
(617) 460-4646

Product Types:
Local Area Networks

Concord Data Systems, Inc.
45 Bartlett St.
Marlborough, MA 01752
(617) 460-0808

Product Types:
Network Management Systems

Consolidated Network Inc.

11701 Borman Drive, Suite 215A
St. Louis, MO 63146
(314) 993-9009

Product Types:
Interstate Switched Facilities
Private Line Facilities

Contel ASC

1801 Research Boulevard
Rockville, MD 20850
(301) 251-8333

Product Types:
Private Line Facilities

Control Data Corp.

8100 34th Avenue South
Minneapolis, MN 55420
(612) 853-8100

Datapoint Corp.

Product Types:
Network Systems Integration
Local Area Networks

Corvus Systems Inc.

160 Great Oaks Boulevard
San Jose, CA 95119
(408) 281-4100

Product Types:
Local Area Networks

CXR Telcom Corp.

755 Ravendale Drive
Mountain View, CA 94043
(415) 965-7100

Product Types:
Network Management Systems

Data General

1 Exchange Plaza
New York, NY 10006
(212) 809-8220

Product Types:
Interstate Switched Facilities
Private Line Facilities

Data General Corp.

4400 Computer Drive
Westboro, MA 01580
(617) 366-8911

Data General Corp.
2155 Leanne Boulevard
Mississauga, ON L5K 2K8
(416) 823-7830

Product Types:
Local Area Networks

Data Switch Corp.

One Enterprise Drive
Shelton, CT 06484
(203) 926-1801

Product Types:
Network Management Systems
Electronic Matrix Switches

Datacomm Management Sciences, Inc.

25 Van Zant Street
East Norwalk, CT 06855
(203) 838-7183

Product Types:
Electronic Matrix Switches
Network Management Systems

Datapoint Corp.

9725 Datapoint Drive
San Antonio, TX 78284
(512) 699-7542

Product Types:
Local Area Networks

Directory of Suppliers

DCA/Cohesive Network

DCA/Cohesive Network

150 Knowles Drive
Los Gatos, CA 95030
(408) 370-4100

Product Types:
T1 Multiplexers

Digilog Inc.

1370 Welsh Road
Montgomeryville, PA 18936
(215) 628-4530

Product Types:
Network Management Systems

Digital Communication Associates

1000 Alderman Drive
Alpharetta, GA 30201-4199
(404) 442-4000

Product Types:
Network Management Systems
Packet Switches

Digital Equipment Corp.

146 Main Street
Maynard, MA 01754-2571
(617) 897-5111

Digital Equipment of Canada, Ltd.
100 Herzberg Road
Kanata, ON K2K 2A6
(613) 592-5111

Product Types:
Network Management Software
Local Area Networks

DMW Commercial Systems, Inc.

202 Hogback Rd.
Ann Arbor, MI 48104
(313) 971-5234

Product Types:
Network Design and Analysis
Communications Management Software
Facilities Management

DNA Networks, Inc.

81 Great Valley Parkway
Malvern, PA 19355
(215) 296-7420

Product Types:
Local Area Networks

Doelz Networks, Inc.

9501 Jeronimo Road
Irvine, CA 92718
(714) 851-2223

Product Types:
Packet Switches

DSC Nestar Systems

1345 Shorebird Way
Mountain View, CA 94043
(415) 969-1777

Product Types:
Local Area Networks

Duquesne Systems

Two Allegheny Center
Pittsburgh, PA 15212
(412) 323-2600; (800) 323-2600; telex
902803

Product Types:
SNA Network Performance Monitoring
Software

Dynatech Communications

991 Annapolis Way
Woodbridge, VA 22191
(703) 550-0011

Product Types:
Network Management Systems
Packet Switches

EDA Instruments Inc.

4 Thorncliffe Park Drive
Toronto, ON M4H 1H1
(416) 425-7800

Product Types:
Packet Switches

Emcom Corp.

101 E. Park Boulevard, Suite 901
Plano, TX 75074
(214) 423-7183

Product Types:
Network Management Systems

Excelan, Inc.

2180 Fortune Drive
San Jose, CA 95131
(408) 434-2300

Product Types:
Local Area Networks

FiberCom Inc.

P.O. Box 11966
Roanoke, VA 24022-1966
(703) 342-6700

Product Types:
Local Area Networks

FiberLAN, Inc.

P.O. Box 12726, 99 T.W. Alexander Drive
Research Triangle Park, NC 27709
(919) 549-6551

Product Types:
Local Area Networks

General DataComm Industries, Inc.

Fibronics International Inc.

Communications Way, Independence Park
Hyannis, MA 02601-1892
(617) 778-0700

Product Types:
Local Area Networks

First Phone

180 Bent Street
Cambridge, MA 02141
(617) 354-5465

Product Types:
Interstate Switched Facilities
Private Line Facilities

FTCC McDonnell Douglas International Telecommunications Co.

90 John Street
New York, NY 10038
(212) 669-9700

Product Types:
Interstate Switched Facilities
Private Line Facilities

Fujitsu America, Inc.

3055 Orchard Drive
San Jose, CA 95134-2017
(408) 946-8777

Fujitsu Canada, Inc.
6280 Northwest Drive
Mississauga, ON L4V 1J7
(416) 673-8666

Product Types:
Network Management Systems

Gandalf Data Inc.

1020 S. Noel Avenue
Wheeling, IL 60090
(312) 541-6060

Gandalf Technologies, Inc.
100 Colonade Road South
Nepean, ON K2E 7J5
(613) 723-6500

Product Types:
Local Area Networks

Gateway Communications, Inc.

2941 Alton Avenue
Irvine, CA 92714
(714) 553-1555

Product Types:
Local Area Networks

General DataComm Industries, Inc.

Route 63
Middlebury, CT 06762-1299
(203) 574-1118

Directory of Suppliers

General DataComm Industries, Inc.

General DataComm Ltd.
2255 Sheppard Avenue East
Willowdale, ON M2J 4Y3
(416) 498-5100

Product Types:
Network Management Systems

GTE

Spacenet Corp.
1700 Old Meadow Road
McLean, VA 22102
(703) 848-1000

Product Types:
Private Line Facilities

Halley Systems, Inc.

2811 Orchard Parkway
San Jose, CA 95134
(408) 434-3500

Product Types:
Local Area Networks

Harris Corp., Business Communications System Div.

P.O. Box 1700
Melbourne, FL 32901
(305) 724-3000

Harris Systems Ltd.
19 Lesmill Road
Don Mills, ON M3B 2T3
(416) 441-2400

Product Types:
Packet Switches

Hewlett-Packard, Information Networks Group

19490 Homestead Rd.
Cupertino, CA 95014
(408) 725-8111; (800) 367-4772

Product Types:
Network Management Systems
Local Area Networks

Hughes LAN Systems

1225 Charleston Road
Mountain View, CA 94043
(415) 966-7400

Product Types:
Hughes LAN Systems
36 Toronto Street, Suite 850
Toronto, ON M5C 2C5
(416) 368-7759

Product Types:
Local Area Networks
Network Access Control Devices

Hughes Network Systems, Inc.

11717 Exploration Lane
Germantown, MD 20874
(301) 428-5500

Product Types:
Packet Switches

Indiana Switch Inc.

P.O. Box 1785
Indianapolis, IN 46206
(305) 869-5200

Product Types:
Interstate Switched Facilities

Industrial Networking, Inc. (INI)

3990 Freedom Circle
Santa Clara, CA 95052
(408) 496-0111

Product Types:
Local Area Networks

The Info Group

46 Park Street
Framingham, MA 01701
(617) 875-7511

Product Types:
Telecom Network Management Systems

Infotext

1067 E. State Parkway
Schaumburg, IL 60173-4559
(312) 490-1155

Product Types:
Telemanagement Systems

Infotron Systems Corp.

9 N. Olney Avenue
Cherry Hill, NJ 08003
(609) 424-9400

Product Types:
Network Management Systems

InteCom Inc.

601 InteCom Drive
Allen, TX 75002
(214) 727-9141
Telecommunications Terminal Systems
2255 Sheppard Avenue East, Suite 440 E
Willowdale, ON M2J 4Y1
(416) 756-4900

Product Types:
Local Area Networks

Integrated Telecom Corp.

630 International Parkway
Richardson, TX 75081
(214) 234-3340

ITT/United States Transmission Systems

Product Types:
T1 network management systems

Intel Corp.

3065 Bowers Avenue
Santa Clara, CA 95052-8065
(408) 987-8080

Product Types:
Local Area Networks

Intelligent Software

4330 Shawnee Mission Parkway
Suite 202
Fiarway, KS 66205
(913) 362-5367

Product Types:
Telemanagement Systems

International Business Machines Corp. (IBM)

Old Orchard Road, Armonk, NY 10504
Contact Your Local IBM Representative

IBM Canada Ltd.
3500 Steeles Avenue East
Markham, ON L3R 2Z1
(416) 474-2111

Product Types:
Network Management Systems
Local Area Networks

International Data Sciences, Inc.

7 Wellington Road
Lincoln, RI 02865
(401) 333-6200

Product Types:
Network Management Systems

IRI Communications

600 Third Avenue
New York, NY 10016
(212) 752-4200

Product Types:
Interstate Switched Facilities
Private Line Facilities

Italcable USA Inc.

6 W. 48th Street
New York, NY 10036
(212) 764-6310

Product Types:
Interstate Switched Facilities
Private Line Facilities

ITT/United States Transmission Systems

100 Plaza Drive
Secaucus, NJ 07096
(201) 330-5000

Directory of Suppliers

ITT/United States Transmission Systems

Product Types:
Interstate Switched Facilities
Private Line Facilities

Kaptronix, Inc.

332 Lincoln Drive
Haworth, NJ 07641
(201) 769-4250

Product Types:
Network Management Software
SNA Software

LDL—Long Distance for Less

C.C.I. Plaza, Suite 211
4561 McDowell Road
Phoenix, AZ 85008
(602) 244-0707

Product Types:
Interstate Switched Facilities
Private Line Facilities

LDX Net, Inc.

15450 S. Outer Forty Road
Chesterfield, MO 63017
(314) 537-8000

Product Types:
Private Line Facilities

Lightnet

600 E. Jefferson Street
Rockville, MD 20852
(301) 738-8100

Product Types:
Private Line Facilities

LiTel Telecommunications Corp.

200 Old Wilson Bridge Road
Worthington, OH 43085
(614) 436-1211

Product Types:
Private Line Facilities

Local Area Telecommunications, Inc.

17 Battery Place, Suite 1935
New York, NY 10004
(212) 509-5111

Product Types:
Private Line Facilities

Make Systems, Inc.

201 San Antonio Circle, Suite 225
Mountain View, CA 94040
(415) 941-9800

Product Types:
Network Management Systems

MCI International, Inc.

2 International Drive
Rye Brook, NY 10573
(914) 937-3444

Product Types:
Network Management
Software and Services
Interstate Switched Facilities
Private Line Facilities

MCI Telecommunications Corp.

1133 19th Street, NW
Washington, DC 20036
(202) 872-1600

Product Types:
Interstate Switched Facilities
Private Line Facilities

Memotec Data, Inc.

600 McCaffrey
Montreal, PQ H4T 1N1
(514) 738-4781

Product Types:
Packet Switches
Protocol Conversion Systems

Memotec Datacom, Inc.

40 High Street
North Andover, MA 01845
(617) 681-0600

Product Types:
Network Management Systems
Packet Switches

Micom Systems, Inc.

4100 Los Angeles Avenue, P.O. Box 8100
Simi Valley, CA 93062-8100
(805) 583-8600

Signatel
195 Riviera Drive
Markham, ON LCR 5J6
(416) 477-9973

Product Types:
Packet Switches
Data PBXs
Network Management Systems

Microtel

7100 W. Camino Real
Boca Raton, FL 33433
(305) 392-2244

Product Types:
Private Line Facilities

MidAmerican

2918 N. 72nd Street
Omaha, NE 68134
(402) 392-6800

NEC America, Inc.

Product Types:
Interstate Switched Facilities
Private Line Facilities

Morino Associates, Inc.

8615 Westwood Center Drive
Vienna, VA 22180-2215
(703) 734-9494

Product Types:
SNA Network Monitoring Software

MOSCOM Corporation

300 Main Street
East Rochester, NY 14445
(716) 385-6440

Product Types:
Telemanagement Software

Motorola Computer Systems Inc.

10700 N. De Anza Boulevard
Cupertino, CA 95014
(408) 255-0900

Product Types:
Local Area Networks

National Telecommunications Network (NTN)

1350 Piccard Drive
Rockville, MD 20850
(310) 258-9717

Product Types:
Private Line Facilities

NBI, Inc.

3450 Mitchell Lane, P.O. Box 9001
Boulder, CO 80301
(303) 444-5710

Product Types:
Local Area Networks

NCR Comten

2700 Snelling Avenue North
St. Paul, MN 55113
(612) 638-7777

NCR/Comten Canada, Ltd.
515 Consumers Rd.
Suite 100
Willowdale, ON M2J 4Z2
(416) 496-1300

Product Types:
Communications Processors
SNA Network Management Systems
Protocol Conversion Systems

NEC America, Inc.

8 Old Sod Farm Road
Melville, NY 11747
(516) 753-7000

Directory of Suppliers

NEC America, Inc.

Product Types:
Integrated Voice Data PBXs

NEC America, Inc.—Data Communications Products

110 Rio Robles
San Jose, CA 95134-1899
(408) 433-1250

Product Types:
Network Management Systems

Netrix Corp.

380 Herndon Parkway
Herndon, VA 22070
(703) 481-0606

Product Types:
Network Management Systems
Packet Switches

Network Dimensions

5339 Prospect Road/Suite 122
San Jose, CA 95129
(408) 446-9598

Product Types:
Network Management Display Software

Network Equipment Technologies, Inc. (N.E.T.)

400 Penobscot Drive
Redwood City, CA 94063
(415) 366-4400

Product Types:
T1 Network Management Systems

Network Management, Inc.

11242 Waples Mill Road
Fairfax, VA 22030
(703) 385-4774

Product Types:
Network Design and Analysis
Network Management Services

Network Systems Corp. (NSC)

7600 Boone Avenue North
Minneapolis, MN 55428
(612) 424-4888

Network Systems, Ltd.
5955 Airport Road, Suite 106
Mississauga, ON L4V 1R9
(416) 676-1663

Product Types:
Local Area Networks

Newbridge Networks, Inc.

13873 Park Center Road, Suite 160
Herndon, VA 22071
(800) 332-1080

Product Types:
T1 Network Management Systems

Nixdorf Computer Corp.

300 Third Avenue
Waltham, MA 02154
(617) 890-3600

Product Types:
Network Management Software

Northern Telecom

2435 N. Central Expressway
Richardson, TX 75080
(214) 301-2128

Product Types:
Local Area Networks
Packet Switches
Network Management Systems

Novell, Inc.

122 E. 1700 South
Provo, UT 84601
(801) 379-5900

Product Types:
Local Area Networks

Novell, Inc., Hardware Products Div.

1610 Berryessa Road
San Jose, CA 95133
(408) 729-6700

Product Types:
Local Area Networks

NYNEX Corp.

400 Westchester Avenue
White Plains, NY 10604
(914) 683-2121

Product Types:
Interstate Switched Facilities
Private Line Facilities

Octocom Systems, Inc.

255 Ballardvale St.
Wilmington, MA 01887
(617) 658-6050

Product Types:
Network Management System

Optimum Electronics, Inc.

425 Washington Avenue, P.O. Box 250
North Haven, CT 06473
(203) 239-6098

Product Types:
Network Management Systems

Prime Computer, Inc.

Oracle, Inc.

20 Davis Drive
Belmont, CA 94002
(415) 598-8000

Product Types:
Network Database Systems

Orchid Technology, Inc.

45365 Northport Loop West
Fremont, CA 94538
(415) 683-0300

Product Types:
Local Area Networks

OST, Inc.

1650 Sycamore Avenue
Bohemia, NY 11716
(516) 563-0400

Product Types:
Network Management Systems

Overseas

Telecommunications Commission (OTC Australia)

535 Fifth Avenue, 24th Floor
New York, NY 10017
(212) 370-3885

Product Types:
Interstate Switched Facilities
Private Line Facilities

Pacific National Telecom

805 Broadway
Vancouver, WA 98668
(206) 696-6969

Product Types:
Private Line Facilities

Pacific Telesis Group

140 New Montgomery Street
San Francisco, CA 94105
(415) 542-9000

Product Types:
Interstate Switched Facilities
Private Line Facilities

Peregrine Systems, Inc.

15707 Rockfield Blvd.
Suite 320
Irvine, CA 92718
(714) 855-3923; fax (714) 855-1538

Product Types:
Data Center Management
Network Database Systems

Prime Computer, Inc.

Prime Park
Natick, MA 01760
(617) 655-8000

Directory of Suppliers

Prime Computer, Inc.

Product Types:
Local Area Networks

Proteon, Inc.

Two Technology Drive
Westborough, MA 01581-5008
(617) 898-2800

Product Types:
Local Area Networks

Pure Data Ltd.

200 W. Beaver Creek Road
Richmond Hill, ON L4N 1B4
(416) 731-6444

Product Types:
Local Area Networks

Quadram Corp.

One Quad Way
Norcross, GA 30093
(404) 923-6666

Product Types:
Local Area Networks

QWEST

17304 Preston Road, Suite 1400
Dallas, TX 75252
(214) 931-3727

Product Types:
Private Line Facilities

Racal-Milgo, Inc.

1601 N. Harrison Parkway
Sunrise, FL 33323
(305) 592-8600

Product Types:
Network Management Systems

Racal-Vadic, Inc.

1525 McCarthy Boulevard
Milpitas, CA 95035
(408) 946-2227

Product Types:
Network Management Systems

Racore Computer Products, Inc.

170 Knowles Drive
Los Gatos, CA 95030
(408) 374-8290

Product Types:
Local Area Networks

Radio Shack, Div. of Tandy Corp.

1700 One Tandy Center
Fort Worth, TX 76102
(817) 390-3011

Product Types:
Local Area Networks

RCA Global Communications

60 Broad Street
New York, NY 10004
(212) 806-7000

Product Types:
Interstate Switched Facilities
Private Line Facilities

RCI Corp.

333 Menlo Park
Rochester, NY 14623
(716) 475-8000

Product Types:
Interstate Switched Facilities
Private Line Facilities

Reliance Communications

4907 Bethesda Avenue, Suite 122
Bethesda, MD 20814
(301) 649-1094

Product Types:
Private Line Facilities

RTP Group

595 East Colorado Blvd/4th floor
Pasadena, CA 91101
(818) 304-9146

Product Types:
Telemanagement Systems

Siecor Corp.

489 Siecor Park
Hickory, NC 28603
(704) 327-5998

Product Types:
Local Area Networks

Siemens Data Systems, Inc.

110 Ricefield Lane
Hauppauge, NY 11788
(516) 435-4000

Product Types:
Packet Switches

Software AG of North America, Inc.

Reston International Center
11800 Sunrise Valley Drive
Suite 917
Reston, VA 22091
(703) 860-5050; (800) 336-3761

Product Types:
Data Center Management
SNA Network Session Management

Stonehouse & Co.

SouthernNet, Inc.

61 Perimeter Park, NE
Atlanta, GA 30341
(404) 458-4927

Product Types:
Interstate Switched Facilities
Private Line Facilities

Southland Fibernet Inc.

6702 Plantation Road
Pensacola, FL 32504
(904) 477-1688

Product Types:
Private Line Facilities

Southwestern Bell

1 Bell Center
St. Louis, MO 63101
(314) 235-9800

Product Types:
Interstate Switched Facilities
Private Line Facilities

Standard Microsystems Corp.

35 Marcus Boulevard
Hauppauge, NY 11788
(516) 273-3100

Product Types:
Local Area Networks

Starnet Corp.

3989 Ruffin Road
San Diego, CA 92123
(619) 569-4022

Product Types:
Interstate Switched Facilities

StarTel

409 N. Texas Avenue
Bryan, TX 77806
(409) 779-2830

Product Types:
Private Line Facilities
Interstate Switched Facilities

STC PLC

10 Maltravers Street
London WC2R 3HA
01 836 8055

Product Types:
Managed Transmission Networks

Stonehouse & Co.

4100 Spring Valley Road
Suite 400
Dallas, TX 75234
(214) 960-1566

Directory of Suppliers

Stonehouse & Co.

Product Types:
Telecommunications Management Software

Stratacom, Inc.

3175 Winchester Boulevard
Campbell, CA 95008
(408) 370-2333

Product Types:
Packet Switches

Symplex Communications Corp.

5 Research Drive
Ann Arbor, MI 48103
(313) 995-1555

Product Types:
Network Management Systems

SynOptics Communications, Inc.

329 N. Bernardo Avenue
Mountain View, CA 94043-5223
(415) 960-1100

Product Types:
Local Area Networks

Systemtech Associates

675 Blinn Court
Los Altos, CA 94022
(415) 344-0670

Product Types:
Network Management Services
Network Management Technology

Technetronic USA Inc.
7927 Jones Branch Drive
Suite 400
McLean, VA 22102
(703) 749-1471

Product Types:
Network Performance and Planning Software

Telco Research

1207 17th Avenue South
Nashville, TN 37212
(615) 329-0031

Product Types:
Telemanagement Software

Telecom Canada
160 Elgin Street
Ottawa, ON K1G 3J4
(613) 567-1081

Product Types:
Interstate Switched Facilities
Private Line Facilities

Telefile

17131 Daimler Street
Irvine, CA 92714
(714) 250-1830

Product Types:
Packet Switches

Telematics International, Inc.

Crown Center, 1415 NW 62nd Street
Fort Lauderdale, FL 33309
(305) 772-3070

Product Types:
Packet Switches

Telenet, A US Sprint Co.

12490 Sunrise Valley Drive
Reston, VA 22096
(703) 689-6000

Product Types:
Packet Switches

Telenex Corp.

13000 Midlantic Drive, P.O. Box 869
Mount Laurel, NJ 08054
(609) 234-7900

Product Types:
Electronic Matrix Switches
Network Management Systems

Teleprocessing Products Inc.

4565 E. Industrial Street, Suite 7-K
Simi Valley, CA 93063
(805) 522-8147

Product Types:
Network Management Systems

Telesphere Network, Inc.

2 Mid America Plaza, Suite 500
Oak Brook Terrace, IL 60181
(312) 954-7700

Product Types:
Interstate Switched Facilities
Private Line Facilities

Tellabs Inc.

4951 Indiana Avenue
Lisle, IL 60532
(312) 969-8800

Product Types:
Packet Switches

Teltone Corp.

10801 120th Avenue NE
Kirkland, WA 98033
(206) 827-9626

Product Types:
Local Area Networks

TRT Telecommunications Corp.

TelWatch Inc.

2905 Wilderness Place
Boulder, CO 80301
(303) 440-4756, (800) 669-1266

Product Types:
T1 Network Management Systems
Telecommunications Management Software

10-NET Communications

7016 Corporate Way
Dayton, OH 45459
(513) 433-2238

Product Types:
Local Area Networks.

3Com Corp.

3165 Kifer Road
Santa Clara, CA 95052-8145
(408) 562-6400

Product Types:
Local Area Networks

Tiara Computer Systems, Inc.

2700 Garcia Avenue
Mountain View, CA 94043
(415) 965-1700

Product Types:
Local Area Networks

Timeplex, Inc.

400 Chestnut Ridge Road
Woodcliff Lake, NJ 07675
(201) 391-1111

Timeplex Canada Inc.
90 Nolan Court, Unit 44
Markham, ON L3R 4L9
(416) 475-1961

Product Types:
Packet Switches
Network Management Systems

Total-Tel USA

140 Little Street
Belleville, NJ 07109
(201) 751-5510

Product Types:
Interstate Switched Facilities
Private Line Facilities

TRT Telecommunications Corp.

1331 Pennsylvania Avenue NW, Suite 1100
Washington, DC 20004
(202) 879-2200

Product Types:
Interstate Switched Facilities
Private Line Facilities

Directory of Suppliers

TRW Information Networks Div.

TRW Information Networks Div.

23800 Hawthorne Boulevard
Torrance, CA 90505
(213) 373-9161

Product Types:
Local Area Networks

Tymnet—McDonnell Douglas Network Systems Co.

2560 N. First Street, P.O. Box 49019
San Jose, CA 95161-9019
(408) 922-0250

Product Types:
Packet Switches
Network Management

US Sprint

2330 Shawnee Mission Parkway
Westwood, KS 66205
(913) 676-3777

Product Types:
Private Line Facilities
Interstate Switched Facilities

U S WEST

7800 E. Orchard Road
Englewood, CO 80111
(303) 793-6500

Product Types:
Interstate Switched Facilities
Private Line Facilities

U S WEST Network Systems, Inc. (USWNSI)

14335 NE 24th Street
Bellevue, WA 98007
(206) 644-8400

Product Types:
VTAM Network Management Interface
SNA Device Inventory System

Ungermann-Bass, Inc.

3900 Freedom Circle
Santa Clara, CA 95052-8030
(408) 496-0111

Ungermann Bass Ltd.
Atria North, Suite W 308
2255 Sheppard Avenue East
Willowdale, ON M2J 2Y1
(416) 494-4426

Product Types:
Local Area Networks

Unisys Corp.

P.O. Box 500
Blue Bell, PA 19424-0001

Product Types:
OSI Management Software
Terminal Network Management Software

Vitalink Communications Corp.

6607 Kaiser Drive
Fremont, CA 94555
(415) 755-6130

Product Types:
Local Area Networks

Vitel International

5 Hanover Square, 12th Floor
New York, NY 10004
(212) 809-9797

Product Types:
Interstate Switched Facilities
Private Line Facilities

VM Software, Inc.

1800 Alexander Bell Drive
Reston, VA 22091
(703) 264-8080

Product Types:
SNA Network Management Interface

Wang Laboratories Inc.

One Industrial Avenue
Lowell, MA 01851
(617) 459-5000

Product Types:
Local Area Networks
Network Management Systems

Waterloo Microsystems Inc.

3597 Parkway Lane
Norcross, GA 30092
(404) 441-9252

Product Types:
Local Area Networks

Zenith Electronics Corp.

Western Digital Corp.

2445 McCabe Way
Irvine, CA 92714
(714) 863-0102

Product Types:
Local Area Networks

Western Union Telegraph Co.

1 Lake Street
Upper Saddle River, NJ 07458
(201) 825-5000

Product Types:
Interstate Switched Facilities
Private Line Facilities

Williams

Telecommunications Co.

1 Williams Center, 45th Floor, P.O. Box 21348
Tulsa, OK 74121
(918) 588-3210

Product Types:
Private Line Facilities

Xerox Corp.

800 Long Ridge Road
Stamford, CT 06904
(Contact Your Local Xerox Representative)

Product Types:
Local Area Networks

Xiox Corporation

577 Airport Blvd.
Suite 700
Burlingame, CA 94010
(415) 375-8188

Product Types:
Telecommunications Management Software

Xtend

171 Madison Avenue/7th floor
New York, NY 10016
(212) 725-2010

Product Types:
Telemanagement Systems

Zenith Electronics Corp.

699 N. Wheeling Road
Mount Prospect, IL 60056
(312) 699-2199

Product Types:
Local Area Networks □

Directory of Communications Consultants

Selecting the right consultant for a particular application involves a great deal of research and evaluation. This annually revised "Directory of Communications Consultants" is designed to help with the critical selection process. Comparison columns list all pertinent information about each firm and the services it provides; we include an entry for "Affiliations," which may help to verify that a prospective consultant is unbiased. The directory contains listings of many voice and data communications consultants. Consultants interested in being included on our mailing list for next year are encouraged to contact us.

Solving communications problems sometimes requires the services of a professional consultant. Choosing the best consultant for a particular application always involves research. This report contains information that can help an organization make this important decision. It examines the idea of consulting in general, communications consulting in particular, and offers some suggestions for the effective use of consultants. Comparison columns listing many firms engaged in such disciplines as telecommunications, data communications, and office automation follow this discussion.

Consultants—Why Use Them?

Consultants help management analyze problems associated with corporate goals, objectives, policies, strategies, administration, and organization. They can recommend solutions to these problems and help implement them if necessary.

Consultants offer a number of benefits, but objectivity is probably the most important. A consultant is generally impartial and free of corporate politics—someone who can get to the heart of an issue without prejudice. A consultant also is more likely to have a wide variety of experiences from which to offer guidance. This can be an important asset while anticipating problems, selecting the best course of action, and predicting the expected results. Analytical skills gained from years of experience often allow the consultant to isolate *real* problems, from *perceived* ones. Finally, one thing consultants generally have, which a manager often does not, is *time*. They can sidestep

the numerous interruptions faced by a manager each day, thus maximizing their skills to benefit the client.

When Should You Use a Consultant?

Before deciding to hire a consultant, an organization must decide whether it actually needs one. This may involve an evaluation of the ability of in-house personnel to complete certain tasks associated with a communications project. The lack of in-house personnel to complete the project is another primary consideration.

Once the need for a consultant has been established, management must decide what the communications consultant should accomplish for the organization. This is essential to conducting a proper search. Additionally, prospective consultants should be evaluated according to the same criteria.

Generally speaking, an organization should consider engaging a consultant's services when one or more of the following conditions exists.

- Money can be saved by upgrading a communications system or by reorganizing a department.
- The in-house staff cannot handle the project.
- An impartial needs assessment must be performed to substantiate the fact that a problem actually exists.
- New ideas or techniques are needed.
- An organization needs help with a problem it has tried, but failed, to solve.
- Specialized skills in such areas as PBX evaluation and network design are needed.
- Industry skills are necessary to help get a project under way and completed on time.
- An impartial opinion on handling system analysis or conversion is required.

Directory of Communications Consultants

Services Available

Communications consultants can provide numerous services, but the client must choose what it needs. The client may require skills in only certain areas, some of which follow.

- Evaluating the current system;
- Taking equipment inventories;
- Verifying service billing records;
- Evaluating and selecting systems, including Request For Proposal (RFP) development when required;
- Integrating operations within the context of overall corporate objectives;
- Identifying the required management information;
- Selecting facilities;
- Designing a network;
- Achieving optimum network performance;
- Controlling costs;
- Installing a system;
- Training on a new communications system; and
- Organizing a communications department.

Costs and Caveats

Understand how a consultant will be paid. Consultants are generally compensated through one of the following means: a per diem or hourly fee; a bracket quote with an upper and lower range; a fixed amount; or a retainer, which helps ensure the consultant's availability over a defined time period. Managers should be particularly wary of a consultant who offers to save the organization money on communications expenses—typically phone bills—for a percentage of the savings. Quite often this individual will not render any other services after making some cost-cutting recommendations. (Fortunately, the number of cost-cutting consultants has diminished greatly.

To avoid problems concerning compensation and the amount and type of services rendered, negotiate a contract with a consultant in which the types of services expected and the amount and means of compensation are clearly defined and agreed upon by

both parties. Doing so can help to avoid many problems that occur “after the fact.”

While consultants generally offer valuable services, they are not infallible; and some in fact, may be problematic. Clients express a number of typical complaints about consultants: they lack experience, they are too young, or they are unable to stick to the client's task. Some clients have complained that a consultant pirated their employees. Others have reported consultants who merely rehashed existing ideas and opinions, offered nothing new, but charged an exorbitant fee; they have also mentioned consultants who are insensitive to their problems. When assessing a consultant's potential, it is a good idea to evaluate how well he or she will fit into the corporate culture. It is also wise to pinpoint any personal traits that could prove troublesome.

Selecting a Consultant

Consultant selection involves several key tasks, which are outlined below.

1. Define all aspects of the assignment.
2. Obtain names of consulting organizations that could meet your requirements. Business associates, communications user groups, the Yellow Pages, and this report are excellent sources.
3. Select at least three candidates for consideration and make reference checks before inviting them for an interview.
4. If, after preliminary discussions, you are satisfied that a candidate could do the job, ask him or her to submit a formal proposal.
5. Study and evaluate the proposals carefully, being sure to use the same criteria for judging each submission.
6. Make your selection.
7. Obtain necessary approvals from senior management to hire the candidate.
8. Negotiate a contract with that consultant.

There are numerous consultants, but it is prudent to narrow the search to a few, good prospects before selection.

The answers to some key questions are essential to narrowing the selection field. These questions include:

Directory of Communications Consultants

- What is the history of the firm and the credentials of its principals?
- What is the consultant's major area of expertise?
- What professional standards does the firm follow?
- Who are the consultant's previous clients?
- How much of the consultant's business is repeat business?
- Can the firm devote the resources necessary to complete a project?
- Does the firm have applicable communications experience?
- Will the firm submit a written proposal stating the objectives and work to be performed?
- Are fees estimated in advance?
- Do fees appear reasonable when compared to those from competing consultants?

After narrowing the selection and thoroughly checking references of applicable candidates, obtain a proposal outlining all aspects of the project. The proposal should cover the following items:

- A definition of the problem.
- A statement of objectives.
- A statement of the scope and nature of the job.
- A discussion of project methodology.
- An organization chart showing duties and responsibilities of each staff member.
- An estimate of the project's duration.
- Financial statements, including a projection of professional fees.

Ensuring a Successful Engagement

The success of a client/consultant relationship generally rests on both sides' willingness to cooperate and

communicate fully. While there are many intricacies involved in conducting a successful project, several key points should be highlighted:

1. Define the project carefully. Make sure you have determined the real problem.
2. Select the consultant carefully. Don't jump at the first one to submit an attractive bid.
3. Establish the responsibilities of both sides. As the client, make a commitment to the project's success. Get assurances that your consultant will do the same.
4. A successful project requires management. Provide effective supervision and support by setting and maintaining the course of action, as well as conducting regular progress reviews.
5. Follow up on the consultant's recommendations. Review the project approximately one month after completion and later at regular intervals, to evaluate further progress and identify problems.
6. Was the project successful? Measure the results and keep management informed.

The Directory

Each listing includes the name, address, and telephone number of the firm; the person to contact for more information; the date founded; the number of active consultants on staff; the geographic area served; and a description of the services, including technology and applications specialties. To help you verify a consultant's objectivity, we have included a section that explores the consultant's affiliations with other companies. The directory is updated annually.

*Once again, we emphasize the importance of **verifying the reputation and objectivity** of the consultant before hiring. Make sure he or she is unbiased and has no conflict of interest or affiliation with a particular supplier.*

Thank You

We extend our thanks to the consulting firms listed in this directory for complying with our requests for information.

Directory of Communications Consultants

FIRM	A-H & Associates Inc.	Able Telecommunications Inc.	Advanced Network Design, Inc.	Alltelsonics
GENERAL INFORMATION				
Address	55 New Montgomery Street Suite 721 San Francisco, CA 94105	515 Kevenaire Drive Milpitas, CA 95035	12391 Lewis Street Garden Grove, CA 92640	6 M Dorado Drive Morristown, NJ 07960
Telephone Number	(415) 974-5335	(408) 945-1484	(714) 750-8811	(201) 993-8265
Date Founded	1985	1984	1978	1982
Number of Active Consultants on Staff	2	3	28	2
Geographic Area Served	Western U.S.	U.S.	U.S. and Canada	Northeastern U.S.
Customer Contact	Al Armand, President	George S. Chow, President	Dave Wiegand President	Patrick Pisaniello, President
TECHNOLOGY SPECIALITIES	Electronic mail; fiber optics; key phone systems; LANs; long-distance alternatives; modems; PBXs; muxes.; network mgt.; office automation; PC comm.; private networks; voice/data integration; voice messaging	Fiber optics; ISDN; LANs; long-dist. alternatives; muxes.; network mgt.; packet swtchnng.; PC comm.; private networks; T1; voice/data integration	Comm. software; electronic mail; fax.; fiber optics; ISDN; LANs; long-dist. alt.; microwave; network mgt.; office auto.; PBXs; PC comm.; pkt. swtchnng.; private networks; SNA; T1; see Comments	Energy mgmt.; fiber optics; ISDN; LANs; microwave; mobile radio; muxes.; network mgt.; pkt. switching; PBXs; private networks; private pay phones; T1; teleconf.; video; voice/data integ.
APPLICATIONS SPECIALITIES	Corp. users; education; financial; government; health care; manufacturing; mult. bldg. environment; telephone companies	Corp. users; education; financial; international; manufacturing; telephone companies	Corporate; education; financial; gov't; health care; international; manufacturing; military; multiple bldg. env.; real estate	Corp. users; education; health care; international; mult. bldg. environment; svc. bureaus; telephone companies
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	No
Strategic Planning	Yes	Yes	Yes	No
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	Yes	No	No
Administrative/Management Services	Yes	No	Yes	Yes
Product Design and Development	No	No	No	Yes
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	No	No
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Consultant; outside cos.
Recommended Software Is Available through:	Outside companies	Outside companies	Consultant; outside cos.	Consultant did not specify
Recommended Services Are Available through:	Outside companies	Consultant	Outside companies	Consultant
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	—	—	Other technologies include teleconferencing; voice/data integration; and voice messaging. Client base is in excess of 1,100 corporations.	Company is branching out to Europe, (prefers Spain and Italy). Company takes advantage of matchmakers seminar offered by Department of Commerce.

Directory of Communications Consultants

FIRM	The Aries Group	ARINC Research Corporation	Janine Asai and Associates	Ken Asten & Associates, Inc.
GENERAL INFORMATION				
Address	1500 Research Blvd. Suite 230 Rockville, MD 20850	2551 Riva Road Annapolis, MD 21401	12861 Palm Street Garden Grove, CA 92640	225 West Broadway, Suite 500 Glendale, CA 91204
Telephone Number	(301) 762-5500	(301) 266-4000	(714) 530-2836	(818) 507-6189
Date Founded	1980	1958	1988	1979
Number of Active Consultants on Staff	8	1,100	17	1
Geographic Area Served	U.S. and Canada	U.S.	U.S., Europe, Japan	U.S.
Customer Contact	Edward T. Garner	Kenneth E. Lyons Director of Marketing	Janine Asai	Ken Asten, President
TECHNOLOGY SPECIALTIES	Private networks	Energy mgt.; LANs; long-dist. alternatives; network mgt.; office automation; private networks; satellite; packet switching	Comm. software; EDI; fax.; key phone systems; muxes.; network mgt.; office automation; T1; PBXs; PC comm.; private networks; private pay phones; teleconf.; voice/ data integ.; voice msg.	Comm. software; electronic mail; fiber optics; key phone sys.; LANs; long dist. alt.; microwave; modems; muxes.; network mgt.; PBXs; private networks; T1; voice msgng; see Comments
APPLICATIONS SPECIALTIES	Corporate users; telephone companies; government agencies	Corporate users; military; utilities; government	Corp. users; education; financial; gov't; health care; international; military; mult. bldg. env.; see Comments	Corp. users; education; financial; government; utilities; manufacturing; mult. bldg. environment telephone cos.
SERVICES PROVIDED				
Equipment Evaluation	No	Yes	Yes	Yes
Training/Education	No	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	No	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	Yes
Administrative/Management Services	No	No	Yes	Yes
Product Design and Development	No	No	Yes	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	No	Yes	Yes
Other	—	Integration analysis and engineering	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Not applicable	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Not applicable	Outside companies	Consultant did not specify	Outside companies
Recommended Services Are Available through:	Consultant	Consultant; outside cos.	Consultant did not specify	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	Consultant did not specify	No
COMMENTS	Specializes in design, planning, and management of voice and data networks. Has designed some of the largest private networks in U.S. and Canada, including 14 for state governments.	An affiliated company (Aeronautical Radio Inc.) is the communications co. supporting the scheduled airlines. In that role company owns, develops, and operates at one of the nation's largest air-ground data link radio services, in addition to an extensive wide area data network supporting airline industry operators.	Other applications are real estate and shared tenant services enterprises.	Other technologies include voice/data integration.

Directory of Communications Consultants

FIRM	Auditel	Bank Data Bank Consultants	BCI Incorporated	A. Bolger & Associates
GENERAL INFORMATION				
Address	109 Lakefront Drive Hunt Valley, MD 21030	Box 80 Carle Place, NY 11514	12200 Sunrise Valley Drive Suite 100 Reston, VA 22091	1740 North Drury Lane Arlington Hgts, IL 60004
Telephone Number	(301) 785-8400	(516) 334-7362	(703) 648-9044	(312) 398-7537
Date Founded	1974	1977	1979	1975
Number of Active Consultants on Staff	5	2	100	2
Geographic Area Served	U.S.	U.S.	U.S.	U.S.
Customer Contact	Karen Finke	Douglas Blaine Kenney, President	Jim Axford, VP Marketing	Alfred E. Bolger, P.E., President
TECHNOLOGY SPECIALITIES	Comm. software; EDI; elec- tronic mail; fax.; fiber optics; interactive a/v; ISDN; key phone systems; LANs; long-distance alternatives; microwave; mobile radio; modems; muxes.; network mgt.; T1; SNA; see Comments	Electronic mail; LANs; microwave; modems; muxes.; network mgt.; office automation; PBXs; PC comm.; T1, private networks; satellite; telex/TWX; video; trading floor systems.	Fiber optics; info. se- curity; ISDN; long-distance alt.; microwave; mobile radio; muxes.; network mgt.; packet switching; PBXs; private networks; satellite; security; SNA; T1; VSAT; voice/data int.	Private networks; voice/data integration; T1; PBXs; telex/TWX; SNA; satellite; muxes.; modems; network mgt.; key phone sys.; long- distance alternatives
APPLICATIONS SPECIALITIES	Corp. users; education; financial; gov't; health care; international; manu- facturing; military; mult. bldg. env.; real estate; See Comments	Corporate users; financial; international	Corporate; education; financial; health care; international; manufact.; telephone cos.; inter- exchange carriers; utils.	Corporate users; manufacturing; shared tenant; phone companies; international; health care; service bureaus
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	No	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	No
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	Yes
Administrative/Management Services	No	No	Yes	No
Product Design and Development	No	Yes	Yes	No
Cost Analysis	Yes	Yes	Yes	No
Competitive Analysis	Yes	No	Yes	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Affiliate; outside cos.; consultant	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Consultant; outside cos.;	Outside companies
Recommended Services Are Available through:	Outside companies	Outside companies	Consultant	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude office automation; pkt. swtchng.; PBXs; PC comm.; private networks; private pay phones; radio paging; satellite; tele- conf.; telex/TWX; video; voice/data integ.; voice messaging; and VSAT. Other applications are svc. bureaus and shared tenant services.	Bank Data Bank special- izes in the communications and computing systems that support banking and Wall Street applications— trading turrets, market information distribution systems, hoot 'n holler communications, back— office computer support, SWIFT/CHIPS, PC LANs, and special hardware/soft- ware for user-designed situations.	All consultants pro- vide have a minimum of ten years background and are active in the indus- try. Firm specifies deliverables in contract.	

Directory of Communications Consultants

FIRM	Bricker and Associates, Inc.	BTM Associates	Wesley Bull & Associates, Inc.	Arthur L. Cader
GENERAL INFORMATION				
Address	625 North Michigan Avenue Suite 1800 Chicago, IL 60611	30 Galway Drive Mendham, NJ 07945	910 Securities Building Seattle, WA 98101	P. O. Box 4308 Sunnyside, NY 11104
Telephone Number	(312) 787-6777	(201) 543-6501	(206) 624-0224	(718) 729-6507
Date Founded	1979	May 1987	1962	1987
Number of Active Consultants on Staff	20	3	6	1
Geographic Area Served	Primarily Midwest, some national	U.S.	International	Continental U.S.
Customer Contact	Paul Himes	Joe Barrett—Principal	Scott K. Church, P.E., V. P., General Manager	Arthur L. Cader, Proprietor
TECHNOLOGY SPECIALTIES	Electronic mail; fax.; info. security; LANs; modems; network mgt.; office automation; PBXs; PC comm.; private networks; voice/data integration; voice messaging	Comm. software; E-mail; energy mgt.; fax; fiber optics; ISDN; key phone systems; LANs; long-dist. alt.; modems; muxes.; net- work mgt.; office auto.; packet switching; PBXs; PC comm.; private networks; T1; see Comments	Fiber optics; ISDN; key phone sys.; LANs; microwave; mobile radio; network mgt.; PBXs; private networks; satellite; security; T1; video; voice/data integration	Key phone systems; network management; PBXs; cost control
APPLICATIONS SPECIALTIES	Corp. users; government; manufacturing; real estate; mult. bldg. environment associations; professional practices	Corp. users; education; financial; mult. building env.; real estate; shared tenant svcs. enterprises; telephone companies	Corp. users; education; gov't; health care; international; military; mult. bldg. env.; shared tenant svc.; phone cos.	Corporate users; financial; multiple building environ- ment; real estate; shared tenant services
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	No
Training/Education	Yes	Yes	No	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	No	No
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	No	No	No	No
Cost Analysis	Yes	No	Yes	Yes
Competitive Analysis	No	Yes	Yes	No
Other	Contract negotiations; implementation; and human resources	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.	Consultant; outside cos.	Consultant; parent company	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	—	No	No	No
COMMENTS	—	Other technologies in- clude teleconf.; voice/ data integ.; and voice messaging. This firm has 24 years of telecom plan- ning, designing and imple- mentation. Its specialty is strategic planning of office systems and telecom services.	Firm has provided pro- fessional engineering and technical assistance services to the tele- communications industry since 1962. It has no manufacturer or supplier affiliation.	—

Directory of Communications Consultants

FIRM	California Telecommunications Inc.	Call Control Systems	Callahan and Associates Inc.	CIMI Corporation
GENERAL INFORMATION				
Address	31360 Via Colinas, Suite 108 Westlake Village, CA 91362	3027 Calle Frontera San Clemente, CA 92672	7 South Bayview Avenue Fairhope, AL 36532	520 Haddon Avenue Haddonfield, NJ 08033
Telephone Number	(818) 991-3211	(714) 492-2838	(205) 928-1559	(609) 354-1088
Date Founded	1973	1983	1976	1982
Number of Active Consultants on Staff	2	3	6	4
Geographic Area Served	California and Southwest	U.S.	Florida, Georgia, Alabama, Mississippi, Louisiana	U.S.
Customer Contact	W. S. Van Derripe, President	Terry J. Smith President	Mark C. Callahan, Tele- communications Consultant	Thomas L. Nolle, President
TECHNOLOGY SPECIALTIES	Communications software; electronic mail; fax.; interactive a/v; ISDN; LANs; long-distance alternatives; modems; muxes.; network mgt.; office automation; packet switchng.; PBXs; PC comm.; SNA; see Comments	Key telephone systems; long-dist. alternatives; network mgt.; PBXs	Communications software; fiber optics; key phone systems; long-distance alternatives; network mgt.; PBXs; private networks; T1; teleconf.; voice messaging; ACD	Comm. software; EDI; electronic mail; fiber optics; info. security; infrared; ISDN; key telephone systems; LANs; long-dist. alt.; modems; muxes.; network mgt.; office auto.; PBXs; see Comments
APPLICATIONS SPECIALTIES	Corp. users; education; financial; gov't; health care; international; manufacturing; military; mult. bldg. env.; real estate; see Comments	Corp. users; education; financial; government; health care; manufac- turing; svc. bureaus; telephone cos.; utilities	Corporate users; education financial; health care; manufacturing; utilities	Corporate; education; financial; health care; international; manufac- turing; mult. bldg. env.; telephone cos.; utilities
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	No	No	No	Yes
RFP Preparation/Evaluation	Yes	No	Yes	No
Strategic Planning	Yes	No	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	Yes
Administrative/Management Services	Yes	No	Yes	No
Product Design and Development	No	No	No	Yes
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	No	Yes	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Consultant; outside cos.
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Consultant; outside cos.
Recommended Services Are Available through:	Outside companies	Consultant; outside cos.	Outside companies	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude private networks; private pay phones; T1; voice/data integra- tion; voice messaging. Other applications are svc. bureaus; shared tenant services; and teleconferencing.	CCS specializes in pro- viding network analysis studies on local service, FX, PL, TL, WATS, 800, calling card/collect & international usage. It also provides ARS/LCR/FRS design services and routing guides.	—	Other technologies in- clude packet switching; PC communications; private networks; satellite; security; SNA; T1; voice/data integration; voice messaging; and VSAT.

Directory of Communications Consultants

FIRM	Combined Telecom Resources	Communication Management	Communication Planning Corporation	Communication Resources
GENERAL INFORMATION				
Address	416 Dunaille Drive Nashville, TN 37217	29 Hones Road Export, PA 15632	8659 Baypine Road Suite 301 Jacksonville, FL 32256	Box 2018 Haddonfield, NJ 08037
Telephone Number	(615) 361-4300	(412) 325-4545	(904) 733-9090	(609) 429-6843
Date Founded	1981	1977	1978	1974
Number of Active Consultants on Staff	2	1	4	2
Geographic Area Served	U.S.	U.S.	U.S.	U.S.
Customer Contact	Mr. John Lovero, President	Bill Ladetue, Owner	R. A. Carter, Sr. Vice President	Not specified
TECHNOLOGY SPECIALITIES	Electronic mail; fax.; fiber optics; ISDN; key phone systems; long-dist. alternatives; microwave; modems; muxes.; network mgt.; office auto- mation; private networks; T1; see Comments	Facsimile; key phone sys- tems; long-dist. alt.; modems; muxes.; network mgt.; private networks; T1; voice/data integration	Comm. software; elec- tronic mail; energy mgt.; fax.; fiber optics; infrared; key phone systems; LANs; microwave; modems; muxes.; network mgt.; PBXs; private networks; video	ISDN; key phone systems; long-dist. alternatives; network mgt.; PBXs; private networks; T1; teleconf.; telex/TWX; voice/data integ.; voice messaging
APPLICATIONS SPECIALITIES	Corp. users; education; financial; gov't; health care; manufacturing; military; mult. bldg. environment	Corp. users; education; manufacturing; mult. bldg. environment	Corp. users; education; financial; gov't; health care; mult. bldg. environment; real estate; shared tenant svc.; telephone cos.	Corp. users; education; financial; health care; manufacturing; mult. bldg. environment
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	No	No	Yes	Yes
RFP Preparation/Evaluation	Yes	No	Yes	Yes
Strategic Planning	Yes	No	Yes	No
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	No
Administrative/Management Services	Yes	No	Yes	Yes
Product Design and Development	No	No	No	Yes
Cost Analysis	Yes	No	Yes	Yes
Competitive Analysis	Yes	No	Yes	No
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.	Outside companies	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude satellite; voice messaging; and VSAT.	Recent projects include the implementation of Megacom and Megacom 800 with FID4 access and DNIS, a voice data network using a combination of 56 KB equipment, statistical multiplex equipment and alternate use equipment for smaller locations. These applications are suited to the medium-to- small networks.	—	The two consultants with this firm have a total of over 70 years experience in telecommunications.

Directory of Communications Consultants

FIRM	Communication Resources Company	Communications Consultants Corporation of Virginia	Communications Network Architects, Inc.	Communications Plus
GENERAL INFORMATION				
Address	17155 Newhope, Suite "N" Fountain Valley, CA 92708	5516 Jessup Road Richmond, VA 23234	P.O. Box 32063 Washington, DC 20007	10408 Antioch Overland Park, KS 66212
Telephone Number	(714) 546-2771	(804) 275-7755	(202) 775-8000	(913) 381-1170
Date Founded	1975	See Comments	1973	1983
Number of Active Consultants on Staff	9	8	12	3
Geographic Area Served	Western U.S.	U.S.	International	Midwest
Customer Contact	J.L. Supernaw, President	L. Thomas Walton	Kathryn A. Dzubeck, Exec. Vice President	Ron Chilcoat President
TECHNOLOGY SPECIALTIES	Electronic mail; fax.; fiber optics; infrared; key phone systems; LANs; long-dist. alternatives; microwave; modems; muxes.; network mgt.; PBXs; private networks; T1; see Comments	Comm. software; fax.; fiber optics; ISDN; key phone systems; LANs; long-dist. alt.; modems; muxes.; network mgt.; PBXs; PC comm.; private networks; T1; see Comments	Comm. software; EDI; electronic mail; fax.; info. security; ISDN; LANs; long-dist. alternatives; modems; muxes.; network mgt.; office automation; pkt. swtchnng.; PBXs; SNA; T1; see Comments	Communications software; infrared; key phone sys- tems; LANs; long-dist. alt.; microwave; modems; muxes.; network mgt.; PBXs; PC comm.; private networks; T1; video; see Comments
APPLICATIONS SPECIALTIES	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. environment	Corp. users; education; financial; gov't; manu- facturing; military; mult. bldg. environment	Corp. users; education; financial; international; manufacturing; mult. bldg. env.; telephone cos.	Corp. users; education; financial; gov't; health care; manufacturing
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	No	Yes	Yes	No
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	No
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	No	No	Yes	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	No	Yes	Yes	Yes
Other	—	Convert job functions into application software	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.; affiliate	Consultant; outside cos.	Outside companies	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude private pay phones; radio paging; voice/data integ.; voice messaging; voice paging. Firm has fully staffed office in Northern California: 505 W. Olive Ave. Suite 120, Sunnyvale, CA 94026 (408) 738-2922.	This firm is a division of Walton & Walton Associates Inc. Two firms were merged: one was founded in 1985; the other in 1960. Other technologies include private pay phones; teleconf.; and voice/ data integration.	Other technologies in- clude PC comm.; private networks; satellite; VSAT; and voice/data integration.	Other technologies in- clude voice/data integra- tion and voice messaging. Network design and analysis, equipment evaluation, and acquisition services are also offered.

Directory of Communications Consultants

FIRM	Communications Technology Associates (CTA)	Computer Communications Systems Engineering Ltd.	COMSUL Ltd.	Consultel, Inc.
GENERAL INFORMATION				
Address	175 Trotta Drive New Milford, NJ 07646	30011 Ivy Glenn Drive, Suite 223 Laguna Niguel, CA 92677	7500 San Felipe Suite 900 Houston, TX 77063	70 Boston Post Road Wayland, MA 01778
Telephone Number	(201) 599-9075	(714) 249-1158	(713) 780-8860	(617) 647-7777
Date Founded	1984	1983	1967	1982
Number of Active Consultants on Staff	5	3	20	5
Geographic Area Served	New York metropolitan area	Western U.S.	International	U.S.
Customer Contact	Peter J. Carroll, President	K. H. Haynes	Michael K. Dillingham	Douglas J. Mitchell, President
TECHNOLOGY SPECIALITIES	EDI; electronic mail; fax.; key phone systems; LANs; long-dist. alt.; modems; muxes.; network mgt.; office auto.; PBXs; pkt. swtchnng.; private networks; SNA; T1; teleconf.; see Comments	Communications software; energy mgmt.; fiber optics; ISDN; key phone systems; LANs; long-dist. alt.; microwave; modems; muxes.; network mgt.; office automation; pkt. swtchnng.; PBXs; PC comm.; VSAT; T1; see Comments	Comm. software; electronic mail; fax.; fiber optics; infrared; interactive a/v; ISDN; LANs; microwave; modems; muxes.; network mgt.; office auto.; see Comments	Electronic mail; fiber optics; infrared; ISDN; key phone systems; LANs; long-dist. alternatives; microwave; modems; muxes.; network mgt.; office automation; PBXs; PC comm.; T1; see Comments
APPLICATIONS SPECIALTIES	Corp. users; education; financial; international; mult. bldg. env.; real estate; svc. bureaus; telephone cos.; utilities	Corp. users; education; financial; gov't; manufacturing; military; real estate; svc. bureaus; shared tenant svcs.; telephone cos.	Corp. users; education; financial; gov't; health care; international; mult. bldg. env.; real est.; see Comments	Corp. users; education; financial; gov't; health care; international; manufacturing; mult. bldg. env.; utilities
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	No	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	Yes	No	Yes
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	Yes	Yes	No	Yes
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	Yes
Other	—	Telemarketing lead generation	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Consultant	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Consultant	Outside companies	Outside companies
Recommended Services Are Available through:	Outside companies	Consultant	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	Consultant did not specify	No
COMMENTS	Other technologies include telex/ TWX; video; voice/data integ.; and voice messaging. This company runs seminars on the securities industry (stocks and bonds) trading.	Other technologies include private networks; private pay phones; satellite; teleconferencing; voice/data integration and voice messaging.	Other technologies include packet switching; PBXs; PC comm.; private networks; radio paging; satellite; SNA; T1; teleconf.; voice/data integ.; voice messaging. Other applications are shared tenant services enterprises.	Other technologies include private networks; satellite; voice/data integration; voice messaging and teleconferencing.

Directory of Communications Consultants

FIRM	Coopers & Lybrand	CyberLink Corporation	Data Comm Connection Inc.	DAVCOM Communications Consultants, Inc.
GENERAL INFORMATION				
Address	1251 Avenue of the Americas New York, NY 10020	1790 30th Street Suite 300 Boulder, CO 80301	14525 Highway 7, Suite 145 Minnetonka, MN 55331	5091 Spring Rock Terrace Roswell, GA 30075
Telephone Number	(212) 536-2000	(303) 447-2122	(612) 936-4010	(404) 993-7770
Date Founded	1898	1982	1983	1983
Number of Active Consultants on Staff	1,500	10	2	2
Geographic Area Served	U.S. and international	U.S.	U.S.	Southeast
Customer Contact	Dennis Conroy, Partner, F. Hall, Director in New York or any U.S. office	Steve Toth Senior Consultant	Jim Egermeier, CEO	David L. Sulhoff, President
TECHNOLOGY SPECIALTIES	EDI; electronic mail; T1; fiber optics; information security; ISDN; LANs; long-dist. alt.; microwave; modems; muxes.; network mgt.; office auto.; PBXs See Comments	Comm. software; energy mgt.; fax.; FCC lnsng. fiber optics; info. security infrared; ISDN; key phone systems; LANs; long-dist. alt.; microwave; modems; muxes.; network mgt.; PBXs; See Comments	EDI; fiber optics; ISDN; LANs; modems; muxes.; network management; packet switching; PC comm.; SNA; private networks; T1; voice/data integration	Comm. software; EDI; ISDN; key phone systems; LANs; long-dist. alt.; microwave; network mgt.; office automation; pkt. swtchnng.; PBXs; PC comm.; private networks; T1; See Comments
APPLICATIONS SPECIALTIES	Corporate users; education; financial; international; manufac- turing; mult. bldg. env.; See Comments	Corporate; education; gov't; military; shared tenant services; multiple building envi- ronment	Corporate users; financial manufacturing	Corporate users; manufacturing; financial
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	No	No	Yes	No
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	No	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	No	No	No
Administrative/Management Services	No	Yes	Yes	No
Product Design and Development	No	Yes	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	No	Yes
Other	—	Engineering and design	Remote site installation and network operations assistance	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Outside companies	Outside companies	Consultant; subsidiary	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude packet switching; PC comm.; private net- works; SNA; telex/TWX; video; VSAT; voice/data integ.; and voice messag- ing. Other applications are real estate; shared tenant services; and telephone companies.	Other technologies in- clude office auto.; packet switching; PC comm.; pri- vate networks; private pay phones; satellite; SNA; T1; voice/data integ.; voice messaging; VSAT; and wire/cable design. This firm is a member of Society of Telecommuni- cations Consultants (STC) and the IEEE.	—	Other technologies in- clude voice/data integ.; and voice messaging.

Directory of Communications Consultants

FIRM	dBrn Associates, Inc.	DCE Communications Consultants Ltd.	DeLong & Associates	Delphi Inc.
GENERAL INFORMATION				
Address	189 Curtis Road Hewlett Neck, NY 11598	201-251 Laurier Ave. West Ottawa, Ontario K1P 5U6	6940 Maycroft Rancho Palos Verde, CA 90274	20 Passaic Ave. Pompton Lakes, NJ 07442
Telephone Number	(516) 569-4557	(613) 563-0091	(213) 377-2426	(201) 839-5770
Date Founded	1982	1975	1983	1976
Number of Active Consultants on Staff	5	2	3	8
Geographic Area Served	Continental U.S.	U.S. and Canada	International	North America and Europe
Customer Contact	Michael F. Finneran, President	Elaine Wade, Consultant	Kenneth DeLong, Managing Partner	Carolyn Park, Office Manager
TECHNOLOGY SPECIALTIES	Comm. software; E-mail; fax.; fiber optics; interactive a/v; ISDN; LANs; long-dist. alt.; microwave; modems; muxes.; network mgt.; office auto.; packet swtchnng.; see Comments	Electronic mail; fax.; FCC licensing; fiber optics; info. security; interactive a/v; ISDN; key phone systems; LANs; modems; muxes.; network mgt.; office automation; pkt. swtchnng.; PBXs; PC comm.; see Comments	EDI; electronic mail; fax.; infrared; ISDN; LANs; long-dist. alternatives; microwave; modems; muxes.; network mgt.; packet swtchnng.; PBXs; PC comm.; private networks; T1; VSAT; see Comments	Communications software; fiber optics; ISDN; LANs; long-dist. alternatives; modems; muxes.; network mgt.; pkt. swtchnng.; PC comm.; SNA; T1
APPLICATIONS SPECIALTIES	Corp. users; education; financial; health care; international; manufactur- ing; mult. bldg. env.; telephone cos.; utilities	Corp. users; education; financial; gov't; health care; manufacturing; military; mult. bldg. env.; telephone cos.; utilities	Corp. users; financial; international; manufactur- ing; mult. bldg. env.; shared tenant services; telephone cos.	Corp. users; financial; gov't; international; manufacturing; military; telephone companies
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	No	No	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	No	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	No	Yes	No
Administrative/Management Services	Yes	No	Yes	No
Product Design and Development	No	No	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	No	Yes	Yes
Other	—	Cable design engineering	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Not applicable	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Not applicable	Outside companies
Recommended Services Are Available through:	Outside companies	Consultant; outside cos.	Not applicable	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude PBXs; PC comm.; private networks; SNA; T1; teleconf.; telex/TWX; video; and voice/data integration. Michael Finneran, the President, is a columnist for Business Communications Review.	Other technologies in- clude security; teleconf.; voice/data integ.; voice messaging; DOC licensing. This firm assists non-Canadian equipment manufacturers with certification from department of communica- tions and type approval (Canadian version of FCC requirements).	Other technologies in- clude teleconf.; telex/ TWX; and voice messaging.	—

Directory of Communications Consultants

FIRM	Delta Telecomm Company (DTC)	The DMR Group	dp Communications Corporation	Enterprise Development International Inc.
GENERAL INFORMATION				
Address	764 West Turner Avenue Roselle, IL 60172	420 Lexington Ave. New York, NY, 10170	120 Broadway New York, NY 10271	5619 Bradley Boulevard Bethesda, MD 20814
Telephone Number	(312) 893-7829	(212) 949-6655	(212) 227-2400	(301) 652-0141
Date Founded	1981	1973	1975	1983
Number of Active Consultants on Staff	3	2000	70	17
Geographic Area Served	International	International. Offices in North America, U.K., and Australia	New York, Washington, DC, and California	International
Customer Contact	Ron Gaj, President	Burnes P. Hollyman, Principal and Telecom. Practice Director	Lou Capolino, EVP	Oliver Dziggel, President
TECHNOLOGY SPECIALTIES	Microwave; fax.; fiber optics; infrared; interactive a/v; key phone systems; LANs; mobile radio; modems; muxes.; network mgt.; office automation; packet switching; PBXs; T1; see Comments	All technologies offered including image applications and technology	Comm. software; fiber optics; interactive a/v; ISDN; key phone systems; LANs; long-dist. alt.; modems; muxes.; network mgt.; office auto.; packet switching; PBXs; PC comm; T1; see Comments	Electronic mail; fax; fiber optics; information security; interactive a/v; ISDN; key telephone sys.; LANs; network mgt.; office automation; packet switchng.; PBXs; private networks see Comments
APPLICATIONS SPECIALTIES	Corp. users; education; international; real estate; manufacturing; mult. bldg. env.; shared tenant services; hotel	All applications	Corporate; financial; health care; international manufacturing, telephone cos.; shared tenant services	Corp. users; education; gov't; international; military; telephone cos.
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	No	Yes	Yes	No
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	No	Yes	Yes	No
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	Yes	Yes	Yes
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	No	Yes	Yes	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	Yes
Other	—	Project management; facilities management; systems integration	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	See Comments	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	See Comments	Consultant; outside cos.	Outside companies
Recommended Services Are Available through:	Consultant	See Comments	Consultant; outside cos.	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude private networks; radio paging; teleconf.; Telex/TWX; video; voice/data integ.; and voice messaging. Member of the Society of Telecommunications Consultants. Elected consultant member to the Telecom. trade Practices Committee of the Better Business Bureau of Chicago and Northern Illinois.	Hardware, software, and services are recommended through the consultant; parent co.; affiliate; subsidiary; and outside companies.	Other technologies include private networks and voice/data integration.	Other technologies in- clude private pay phones; radio paging; satellite; security; teleconf.; telex/TWX; video; voice/ data integration; voice messaging; and VSAT.

Directory of Communications Consultants

FIRM	The Ergotec Group	Farradyne Systems Inc.	Florida Phone Consultants	FMS Telecommunications Inc.
GENERAL INFORMATION				
Address	559 Bloomfield Avenue Montclair, NJ 07042	3206 Monroe Street Rockville, MD 20852	3601-C South Conway Road Orlando, FL 32814	262 Cedar Street Ashland, MA 01721
Telephone Number	(201) 509-9000	(301) 468-5568	(407) 282-2053	(508) 881-4478
Date Founded	1984	1984	1984	1983
Number of Active Consultants on Staff	4	24	2	4
Geographic Area Served	International	U.S. and Canada	Eastern U.S.	Northeast
Customer Contact	Daniel I. Stusser, President	John F. (Ted) Pugh, Vice President	Wayne F. Freeman, President	Edward R. Rivernider, Manager
TECHNOLOGY SPECIALTIES	Communications software; fax; key telephone systems; network mgt.; PBXs; telemanagement systems and services; EDI	Comm. software; EDI; elec- tronic mail; fax; fiber optics; info. security; infrared; interactive a/v; ISDN; key phone systems; LANs; long-dist. alt.; microwave; mobile radio; modems; muxes; network mgt.; see Comments	Comm. software; electronic mail; fax; fiber optics; infrared; ISDN; key phone systems; LANs; long-dist. alt.; microwave; modems; muxes; network mgt.; packet swtchnng.; PBXs; PC comm.; T1; see Comments	Electronic mail; fiber optics; infrared; ISDN; key phone systems; LANs; long-dist. alternatives; microwave; modems; muxes.; network mgt.; pkt. swtchnng.; PBXs; PC comm.; private networks; private pay phones; see Comments
APPLICATIONS SPECIALTIES	Corp. users; government; military; mult. building env.; real estate; service bureaus; shared tenant	Corp. users; gov't; health care; military; mult. bldg. env.; real estate; shared tenant svc.; utilities; telemetry and realtime control	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. env.; real estate; shared tenant svcs.; util- ities; telephone cos.	Corp. users; education; financial; gov't; health care; mult. bldg. env.; service bureaus; shared tenant; telephone cos.; utilities
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	No	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	No	Yes	Yes	Yes
Network Planning and Design	No	Yes	Yes	Yes
Market Research and Analysis	Yes	Yes	No	No
Administrative/Management Services	No	Yes	Yes	Yes
Product Design and Development	No	Yes	No	No
Cost Analysis	No	Yes	Yes	Yes
Competitive Analysis	No	Yes	No	Yes
Other	—	Software development	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Consultant did not specify
Recommended Software Is Available through:	Outside companies	Consultant; outside cos.	Outside companies	Consultant did not specify
Recommended Services Are Available through:	Consultant; affiliates	Consultant; outside cos.	Outside companies	Consultant did not specify
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	—	Other technologies include office automation; pkt. swtchnng.; PBXs; PC comm.; private networks; satellite; SNA networks; T1; teleconf.; telex/TWX; video; voice/data integ.; voice messaging. Farradyne Systems performs testing of existing systems, installation management, systems, installation mgt., and system specifications. It also offers network modeling capabilities.	Other technologies include private networks; private pay phones; satellite; teleconferencing; voice/data integration; voice messaging; and VSAT.	Other technologies are: radio paging; satellite; SNA networks; T1; voice/data integration; voice messaging; and VSAT.

Directory of Communications Consultants

FIRM	General Network Corporation	Grove Associates Inc.	Henkels & McCoy, Inc. Engineering Division	Harvey S. Hershkowitz Associates, Inc.
GENERAL INFORMATION				
Address	25 Science Park New Haven, CT 06511	18 B Lenox Pointe Atlanta, GA 30324	Jolly Road Blue Bell, PA 19422	519 Bloomfield Avenue Suite 4D Caldwell, NJ 07006
Telephone Number	(203) 786-5140	(404) 233-9611	(215) 283-7763	(201) 226-2059
Date Founded	1984	1981	1928	1984
Number of Active Consultants on Staff	10	4	20	1
Geographic Area Served	U.S. and international	Southeastern U.S.	U.S.	U.S.
Customer Contact	Adrian Zainwel, V.P. Sales Dr. Jason Liu, President	George Grove, President Peter de Golian, Consultant	J. A. Mulhern, manager Comm. Consulting	Harvey S. Hershkowitz, President
TECHNOLOGY SPECIALTIES	Comm. software; electronic mail; facsimile; ISDN; key phone systems; LANs; long-dist. alternatives; modems; muxes; network mgt.; office automation; pkt. swtchng.; PBXs; PC comm.; private networks; VSAT; see Comments	Fax.; fiber optics; infrared; ISDN; key phone systems; long-dist. alt.; network mgt.; office automation; PBXs; PC comm.; private networks; T1; teleconferencing; voice/data integration; voice messaging	Energy mgmt.; fiber optics; ISDN; key phone systems; LANs; long-dist. alternatives; microwave; modems; muxes; network mgt.; PBXs; PC comm.; private networks; SNA; T1; voice/data integ.; VSAT; complex wiring	EDI; electronic mail; PBX; fiber optics; infrared; ISDN; LANs; long-distance alt.; modems; muxes.; network mgt.; packet switch; ing; PC comm.; private networks; T1; voice/data integration
APPLICATIONS SPECIALTIES	All applications	Corp. users; financial; gov't; health care; manufacturing; mult. bldg. env.; real estate; shared tenant services	Corp. users; education; financial; gov't; health care; mult. bldg. env.; shared tenant svcs.; telephone cos.; utilities	Corporate; financial; telephone companies
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	No	Yes
Administrative/Management Services	Yes	Yes	No	No
Product Design and Development	Yes	No	No	Yes
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	No	No	Yes
Other	—	—	Installation/implementation project management	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Consultant; outside cos.	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Consultant; outside cos.	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.	Outside companies	Consultant; outside cos.	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies include voice/data integration. General Network Corporation has developed a family of proprietary software and tools for telecom./computer network modeling analysis, optimization and management.	—	Over 65 years experience in all facets of communications systems and solutions.	This firm has had over 25 years in the communications field including end-user, vendor, and consultant experience.

Directory of Communications Consultants

FIRM	E.C. Hunter Associates, Inc.	Integrated Software Resources, Inc.	The International Project Management Group, Inc	International System Strategies, Inc.
GENERAL INFORMATION				
Address	132 Atkinson Avenue Syracuse, NY 13207	5728 Major Boulevard Suite 500 Orlando, FL 32819	P.O. Box 3606 Glyndon, MD 21071	Building 22, 1640 Powers Ferry Road Atlanta, GA 30067
Telephone Number	(315) 476-3811	(407) 351-9331	(301) 833-9440	(404) 980-2595
Date Founded	1979	1981	1980	1988
Number of Active Consultants on Staff	4	15	14	1
Geographic Area Served	U.S. and Canada	International	Mid-Atlantic	Southeast
Customer Contact	Everest C. Hunter, President	R. Lemuel Lasher, Vice President Marketing	Chuck Gibson	Clayton Bell, Senior Consultant
TECHNOLOGY SPECIALTIES	Comm. software; EDI; E-mail; fax; fiber optics; info. security; infrared; interactive a/v; ISDN; key phone systems; LANs; long-dist. alt.; microwave; mobile radio; modems; see Comments	Comm. software; EDI; network mgt.; pkt. switching; private networks; WANs; protocol conversion message switching & conversion; network design & integration	Fiber optics; key phone systems; long-distance alternatives; network mgt; PBXs; private pay phones; building wiring	Communications software; ISDN; LANs; modems; muxes.; network mgt.; pkt. swtchng.; PC comm.; private networks; security; SNA networks
APPLICATIONS SPECIALTIES	Corp. users; education; financial; gov't; manufacturing; shared tenant svc.; telephone cos.; utilities	Corp. users; financial; gov't; international; svc. bureaus; telephone cos.; airlines/airports	Education; gov't; mult. bldg. env.; real estate; shared tenant services; architects and engineers	Corp. users; education; gov't; international; service bureaus
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	No	Yes
RFP Preparation/Evaluation	Yes	Yes	No	Yes
Strategic Planning	Yes	Yes	No	Yes
Network Planning and Design	Yes	Yes	No	Yes
Market Research and Analysis	Yes	Yes	Yes	Yes
Administrative/Management Services	Yes	No	Yes	No
Product Design and Development	No	Yes	No	No
Cost Analysis	Yes	No	Yes	No
Competitive Analysis	Yes	Yes	Yes	No
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Consultant; affiliate; outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Consultant; affiliate; outside companies	Outside companies	Consultant; outside cos.
Recommended Services Are Available through:	Outside companies	Consultant; outside cos.	Outside companies	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	Yes; sometimes consultant receives hardware royalty from manufacturer	No	No
COMMENTS	Other technologies include muxes.; network mgt.; office auto.; pkt. swtchng.; PBXs; PC comm.; private networks; private pay phones; radio paging; satellite; security; T1; teleconf., voice/data integ.; voice messaging; VSAT; and wire and cable technologies cost reduction.	—	—	—

Directory of Communications Consultants

FIRM	Jamison & Associates	JIM Associates Inc.	Richard N. Kaufman & Associates Ltd.	Kel-Mat Corporation
GENERAL INFORMATION				
Address	9204 LaRiviera Drive Sacramento, CA 95826	543 West Peachtree Street Suite 2620 Atlanta, GA 30308	495 Connecticut Ave. Norwalk, CT 06854	P.O. Box 17117 Pittsburgh, PA 15235
Telephone Number	(916) 361-3332	(404) 881-8590	(203) 838-2000	(412) 243-9370
Date Founded	1983	1970	May 1974	1978
Number of Active Consultants on Staff	5	3	5	10
Geographic Area Served	California and 11 western states	U.S. and Canada	U.S. and international	U.S and Canada
Customer Contact	Don L. Jamison or Jeanne Jamison	Joseph Massey, President; Glenn Thompson, Consultant; David Douglass, Consultant	Richard Kaufman, President	Daniel A. Kelley, VP Engineering
TECHNOLOGY SPECIALITIES	Comm. software; electronic mail; energy mgt. fax.; info. security; ISDN; key phone systems; LANs; modems; network mgt.; office automation; PBXs; PC comm.; security; SNA; see Comments	Comm. software; electronic mail; fax; fiber optics; infrared; interactive a/v; ISDN; key phone systems; LANs; long-dist. alternatives; microwave; modems; muxes; network mgt.; office automation; pkt. swtchng.; PBXs; see Comments	Comm. software; electronic mail; fax; fiber optics; key phone systems; LANs; long-dist. alternatives; microwave; mobile radio; modems; muxes; network mgt.; office automation; PBXs; PC comm.; private nets.; see Comments	Comm. software; EDI; electronic mail; energy mgt.; fiber optics; information security; infrared; interactive a/v; LANs; microwave; muxes.; network mgt.; office automation; PC comm.; see Comments
APPLICATIONS SPECIALTIES	Corp. users; government; health care; manufacturing; military; utilities; mult. bldg. env.; shared tenant svcs.	Corp. users; education; financial; health care; mult. bldg. environment; shared tenant svcs.; telephone cos.; utilities	Corp. users; financial; gov't; health care; international; manufacturing; mult. bldg. env.; svc. bureaus; telephone companies; utilities	Corp. users; financial; health care; manufacturing; multiple building environment
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	Yes	No	No
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	Yes	Yes	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	Yes
Other	—	—	—	Turnkey installation products and mgt.
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Consultant; outside cos. parent co.; affiliate
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Consultant
Recommended Services Are Available through:	Outside companies	Outside companies	Outside companies	Consultant
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies include private networks; voice/data integration; and voice messaging.	Other technologies include PC comm.; private networks; private pay phones; radio paging; T1; voice/data integ.; voice messaging; and VSAT.	Other technologies include radio paging; satellite; T1; teleconf.; telex/TWX; voice/data integ.; voice messaging; and VSAT. Current area of focus is international communications for third world companies.	Other technologies include satellite; security; T1; video; voice/data integ. This firm provides total premise integration of voice, video, and data, using multimedia transportation methods.

Directory of Communications Consultants

FIRM	Knight, O'Connor and Associates, Inc.	LTS Data Communications, Inc.	D. L. Marshall & Associates, Inc.	MBG Associates, Ltd.
GENERAL INFORMATION				
Address	2500 City West Boulevard Suite 2150 Houston, TX 77042	43 Roydon Drive North Merrick, NY 11566	5750 Fort Caroline Road Jacksonville, FL 32211	370 Lexington Ave. Suite 2008 New York, NY 10017
Telephone Number	(713) 780-3226	(516) 783-0111	(904) 743-1111	(212) 687-8580
Date Founded	1984	May 1986	1974	1977
Number of Active Consultants on Staff	6	4	3	6
Geographic Area Served	Gulf Coast and Central Midwest, some U.S. and int.	International	Southeastern U.S.	U.S.
Customer Contact	Bob Falnders, Manager	Jack Gencarelli, President	Don Marshall President	Michael B. Greenspan, President
TECHNOLOGY SPECIALTIES	Comm. software; energy mgt.; fiber optics; information security; key phone sys.; LANs; long-dist. alt.; modems; network mgt.; PBXs; office auto.; PC comm.; private networks; security; T1; teleconf.; telex/TWX; video; see Comments	Comm. software; EDI; Elec- tronic mail; fax.; FCC licensing; fiber optics; info. security; ISDN; key phone systems; LANs; long-dist. alt.; microwave; modems; muxes.; network mgt.; pkt. swtchnng.; PBXs; SNA; see Comments	Comm. software; elec- tronic mail; ISDN; key phone systems; long-dist. alt.; network mgt.; PBXs; private networks; private pay phones; T1; voice/ data integration; voice messaging	Communications software; long-distance alter- natives; network management; private networks
APPLICATIONS SPECIALTIES	Corp. users; financial; gov't; health care; international; manufactur- ing; military; mult. bldg. env.; real estate	Corp. users; financial; international; mult. bldg. env.; svc. bureaus; shared tenant svcs.	Corporate users; financial; government; health care; multiple building environment	Corporate users; financial; service bureaus
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	No
Training/Education	Yes	No	No	No
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	No	Yes	No
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	No	No	No
Administrative/Management Services	Yes	No	Yes	Yes
Product Design and Development	Yes	No	No	Yes
Cost Analysis	Yes	No	No	Yes
Competitive Analysis	Yes	No	Yes	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Consultant; outside cos.	Outside companies	Outside companies	Consultant; outside cos.
Recommended Services Are Available through:	Outside companies	Consultant; outside cos.	Outside companies	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude voice/data integ- ration; voice messaging; and data acquisition and processing.	Other technologies in- clude private networks; T1; telex/TWX; and voice/data integ. Firm provides con- tracts help to project or facility mgt. Firm's con- cept is to provide a combined service to customers to complete a project or to update a service for them, while they manage their function internally within their company.	Don Marshall, principal consultant, has been a consultant since 1974 and President of the Society of Telecommunications dur- ing 1985. Majority of firms clients are large multilocation corporations with systems ranging from 150 to 4,000 stations.	—

Directory of Communications Consultants

FIRM	Belden Menkus, Consultant	Mid-Com Consultants Inc.	Patrick M. Monahan & Associates Inc.	National Consulting Systems, Inc.
GENERAL INFORMATION				
Address	Box 129 Hillsboro, TN 37342	16801 Euclid Avenue Cleveland, OH 44067	907 North Elm Street Hinsdale, IL 60521	9910 North 48th Street, Omaha, NE 68152
Telephone Number	(615) 728-2421	(216) 531-1911	(800) 747-2211 (312) 920-5380	(402) 453-9292
Date Founded	1969	1967	1975	1984
Number of Active Consultants on Staff	2	5	11	9
Geographic Area Served	International	U.S.	U.S.	U.S.
Customer Contact	Not specified	Charles Ault, Marketing Manager	Ronald M. Schwertfeger, Account Executive	James Beatty, President
TECHNOLOGY SPECIALTIES	Comm. software; EDI; electronic mail; fiber optics; info. security; interactive A/V; ISDN; key phone systems; LANs; long-dist. alt.; microwave; network mgt.; office auto.; PBXs; PC comm; private networks; see Comments	Electronic mail; fax; key phone systems; long-dist. alternatives; modems; muxes; network mgt.; PBXs; private networks; T1; voice/data integration; voice messaging	All technologies	Key phone systems; LANs; long-dist. alt.; modems; network mgt.; office automation; PBXs; PC comm.; teleconf.; voice/data integration; voice messaging
APPLICATIONS SPECIALTIES	Corporate; financial; international; service bureaus; telephone companies	Corporate users; education; financial; gov't; health care; manufacturing; multiple building environment; real estate	All applications	Corporate users; education; health care; telephone companies
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	No	Yes	Yes
RFP Preparation/Evaluation	No	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	No	Yes	Yes	Yes
Market Research and Analysis	Yes	No	Yes	Yes
Administrative/Management Services	Yes	Yes	Yes	No
Product Design and Development	Yes	No	Yes	Yes
Cost Analysis	No	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	No
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Not applicable	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Not applicable	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant	Outside companies	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies include private pay phones; security; teleconferencing; and voice messaging.	—	Additional office: 250 W. Coventry Court, Milwaukee, WI 53217, (414) 351-8955, (800) 747-2211.	NCS also handles telephone bills for its clients and universities in developing telecommunications.

Directory of Communications Consultants

FIRM	Network Design and Analysis Corporation (NDA)	Network Resources Inc.	Omnicom, Inc.	Omnicom, Inc.
GENERAL INFORMATION				
Address	60 Gough Road, Unit 2A Markham, Ontario L3R 8X7	210 N. Bassett Street Madison, WI 53703	115 Park Street, SE Vienna, VA 22180	325 John Knox Road Ste E-204 Tallahassee, FL 32303
Telephone Number	(416) 477-9534	(608) 258-7050	(703) 281-1135	(904) 386-3180
Date Founded	January 1983	1987	1982	1982
Number of Active Consultants on Staff	4	2	7	12
Geographic Area Served	International	U.S.	International	International
Customer Contact	Baris Dortok, President	Not specified	Paul Luebbe Technical Director	—
TECHNOLOGY SPECIALTIES	Private networks; comm. software; network management; SNA; packet switching; T1	LANs; long-distance alternatives; network management; PBXs; private networks; voice/data integration	Comm. software; EDI; electronic mail; fax.; information security; ISDN; LANs; network mgt. PC comm.; security; open systems inter- connection	Comm. software; EDI; electronic mail; FCC licensing; interactive a/v; ISDN; key phone systems; LANs; microwave; mobile radio; modems; muxes.; network mgt.; PBXs; see Comments
APPLICATIONS SPECIALTIES	Corporate users; government agencies; financial; international; telephone operating companies	Corp. users; education; gov't; information technology vendors	Corp. users; government; international; manu- facturing; military; telephone companies	Corporate users; government
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	No
Administrative/Management Services	No	Yes	Yes	Yes
Product Design and Development	Yes	No	Yes	No
Cost Analysis	Yes	Yes	No	Yes
Competitive Analysis	Yes	Yes	No	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Consultant	Outside companies	Consultant; outside companies	Outside companies
Recommended Services Are Available through:	Consultant	Consultant; outside cos.	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	NDA provides personal computer-based network design, management and performance analysis pro- ducts coupled with the consulting services.	—		Other technologies in- clude packet switching; PC comm.; private net- works; radio paging; SNA networks; T1; video; voice/data inte- gration; and voice messaging.

Directory of Communications Consultants

FIRM	OptiCom, Inc.	OTM Engineering, Inc.	Pacific Netcom, Inc.	Peat Marwick Main & Co.
GENERAL INFORMATION				
Address	9616 Woodstream Drive Richmond, VA 23233-2907	3101 Bee Caves Road, Suite 325 Austin, Texas 78746	12470 S.W. First Street P.O. Box 1957 Beaverton, Oregon 97005	P.O. Box C-96024 Bellevue, WA 98009
Telephone Number	(804) 747-8646	(512) 328-8801	(503) 641-3933	(206) 455-0111
Date Founded	1981	1983	1983	1898
Number of Active Consultants on Staff	6	3 plus 2 Associates	7	55 telecom; 1800 total
Geographic Area Served	Continental U.S.	Southeast, Southwest and international	The Pacific Northwest	International
Customer Contact	Raymond C. Leffer, President	James Sinopoli, President, P. E.	D. J. Dougherty Marketing Director	See Comments
TECHNOLOGY SPECIALTIES	Comm. software; elec- tronic mail; fax.; ISDN; infrared; key phone systems; LANs; long-distance alternatives; microwave; modems; muxes.; network mgt.; office automation; SNA; see Comments	Electronic mail; fiber optics; interactive a/v; ISDN; key phone systems; LANs; long-dist. alt.; microwave; network mgt.; pkt. swtchnng.; PBXs; T1; teleconf.; private nets.; video; see Comments	Comm. software; elec- tronic mail; fax.; key phone systems; LANs; long-dist. alt.; micro- wave; modems; muxes.; network mgt.; office autoamtion; packet switch- ing; see Comments	All technologies, including "E-911"
APPLICATIONS SPECIALTIES	Corp. users; education; financial; gov't; health care; international; manu- facturing; mult. bldg. env.; see Comments	Corp. users; education; financial; gov't; inter- national; manufact.; military; real estate; see Comments	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. env.; shared tenant; utilities	Corp. users; education; financial; gov't; health care; manufacturing; multiple building env.; real estate; see Comments
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	Yes
Administrative/Management Services	Yes	Yes	Yes	No
Product Design and Development	No	No	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	Yes
Other	—	—	—	Needs analysis
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude office automation; pkt. swtchnng.; PBXs; PC comm.; private networks; private pay phones; radio paging; satellite; T1; teleconf.; telex/TWXX; voice/data integ.; and voice messaging. Other applications are real estate; svc. bureaus; shared tenant services; and telephone companies.	Other technologies in- clude radio paging; voice/data integ.; and voice messaging. Other applications include multiple building environment; shared tenant services; and telephone companies.	Other technologies in- clude PBXs; PC comm.; private networks; private pay phones; radio paging; T1; video; voice/data integration; voice messaging.	Contact on East Coast is Joe Quiggley at (215) 299-3100. Contact in Midwest is Tony Lake at (314) 444-1400. Contact on West Coast is Gary Boyd at (206) 455-0111. Other applications are telephone companies and utilities.

Directory of Communications Consultants

FIRM	Performance Navigation, Inc.	Pho-net-ex, Inc.	Pyramid Communications	RAK Associates
GENERAL INFORMATION				
Address	28 Summit Avenue Hackensack, NJ 07601-1263	1000 Jorie Boulevard Suite 110 Oak Brook, IL 60521	P.O. Box 15288 Sacramento, CA 95851	17894 Clifton Park Lane Cleveland, OH 44107
Telephone Number	(201) 487-0881	(312) 990-2000	(916) 447-1900	(216) 228-2045
Date Founded	1985	1979	1979	1965
Number of Active Consultants on Staff	5	2	3	1
Geographic Area Served	U.S. and Europe	U.S.	U.S.	U.S.
Customer Contact	Mr. Kornel Terplan, President	Jay C. Wheeler, President	Bob Frank, Chief Consultant	Richard A. Kuelin
TECHNOLOGY SPECIALTIES	Communications software; ISDN; network mgt.; pri- vate networks; SNA; T1	Electronic mail; fax.; ISDN; key phone sys.; LANs; long-dist. alt.; microwave; muxes.; T1; network mgt.; PBXs; PC comm; office auto; voice mssng.; private networks; see Comments	Electronic mail; fax.; interactive a/v; key phone systems; LANs; long-dist. alternatives; mobile radio; modems; network mgt.; office automation; PBXs; PC communications; see Comments	Fax; fiber optics; info. security; infrared; ISDN; key phone systems; LANs; long-dist. alt.; microwave; modems; T1; muxes.; network mgt.; pkt. swtchnng.; PBXs; PC comm.; see Comments
APPLICATIONS SPECIALTIES	Financial; international; manufacturing; service bureaus; telephone operating companies	Corporate users; multiple building environments; financial institutions; health care facilities; government; see Comments	Corp. users; education; financial; gov't; health care; svc. bureaus; shared tenant svcs.; telephone cos.; utilities	All applications
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	No	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	No	Yes
Administrative/Management Services	No	Yes	Yes	Yes
Product Design and Development	Yes	No	No	Yes
Cost Analysis	No	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	Yes
Other	—	—	Recovery of communica- tions taxes paid in error by exempt organizations	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.	Consultant; outside cos.	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Major emphasis is on multinational cor- porations.	Other technologies in- clude private pay phones; teleconf.; telex/TWX; and voice/data inte- gration. Other app- lications include education; manufacturing; real estate; and service bureaus.	Other technologies in- clude private networks; private pay phones; radio paging; telex/TWX; voice messaging; communications taxes. Certain groups are exempt from federal taxes or communications services. Firm files exemptions certificates to stop the future billing of the tax and files claims to recover taxes paid in error the past 3 years.	Other technologies in- clude private networks; private pay phones; satellite; security; teleconf.; telex/TWX; voice messaging; and VSAT.

Directory of Communications Consultants

FIRM	RAM Communications Consultants Inc.	Murray H Robinson & Associates Ltd.	Michael L. Rothberg Associates	The RTP Group, Inc.
GENERAL INFORMATION				
Address	1152 St. George Avenue Avenel, NJ 07001	P.O. Box 925 Manotick, Ontario K0A 2N0	27 Heather Drive Somerset, NJ 08873	595 E. Colorado Blvd. Suite 411 Pasadena, CA 91101
Telephone Number	(201) 636-6970	(613) 692-4780	(201) 247-0377	(818) 304-9146
Date Founded	1983	1974	1980	November 22, 1982
Number of Active Consultants on Staff	15	3	6	9
Geographic Area Served	Continental U.S. and Hawaii	Canada, Western Europe, U.K., Northeast U.S.	International	Continental U.S.
Customer Contact	Michael W. Hunter, President	—	Judy Morgenstern, VP	Mr. Gary I. DeLong Vice President
TECHNOLOGY SPECIALITIES	Energy mgt.; FCC licensing; fiber optics; infrared; ISDN; key telephone systems; LANs; microwave; mobile radio; network mgt.; packet switching; PBXs; PC comm.; private networks; radio paging; T1	EDI; electronic mail; facsimile; fiber optics; info. security; ISDN; LANs; long-distance alt.; modems; muxes.; network mgt.; packet switching; PBXs; PC comm.; SNA; see Comments	See Comments	T1; long-distance alternatives; voice/data integration; comm. soft- ware; facsimile; ISDN; key phone systems; PBXs; network management; private networks; voice messaging
APPLICATIONS SPECIALTIES	Education; government; health care; utilities	Corporate users; financial institutions; government agencies; international education; manufacturing; multiple building environment; service bureaus	Corp. users; education; financial inst.; government agencies; health care; international; manufacturing facilities; military; mult. building environment	Corporate; health care; education; financial; gov't; manufacturing; mult. bldg. environment; see Comments
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	Yes	Yes	Yes
Administrative/Management Services	Yes	Yes	Yes	No
Product Design and Development	No	No	Yes	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	No	Yes	Yes	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Consultant; outside cos.	Outside companies	Consultant; affiliate; outside companies	Consultant; affiliate
Recommended Services Are Available through:	Outside cos.	Outside companies	Consultant; affiliate; outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Branch offices located in Boston, Detroit, Seattle, and Honolulu.	Other technologies in- clude private networks; satellite; security; T1; teleconf.; voice/data integ.; voice messaging; and VSAT. Member, Society of Telecommunications Consultants since 1976.	Comm. software; EDI; E-mail; fax; fiber op- tics; info. security; infrared; interactive a/v; ISDN; LANs; modems; muxes; net. mgt.; ofc. aut.; packet switching; PBX; PC comm; private networks; SNA; T1; voice/data integration; and voice messaging.	Other applications in- clude shared tenant svcs. and some work (market research/training) done for other common carriers. RTP special- izes in voice network engineering and developing and marketing a full library of PC-based telecom soft- ware, including call ac- counting and telephone bill management packages.

Directory of Communications Consultants

FIRM	J.R. Schneider & Associates	SCS Engineering Inc.	Sky Wave Associates Ltd.	Southern Telecommunications Consultants
GENERAL INFORMATION				
Address	9861 E. Doubletree Ranch Rd. Scottsdale, AZ 85258	9603 Barcelona San Antonio, TX 78230	3 Country Club Court Mt. Sinai, NY 11766	4963 Meadowbrook Road Birmingham, AL 35242
Telephone Number	(602) 860-0040	(512) 340-7622 (512) 349-7842	(516) 474-2017	(205) 991-5249
Date Founded	December 1, 1965	1983	1982	1984
Number of Active Consultants on Staff	3	5	7	1
Geographic Area Served	U.S., Canada, Europe	U.S.	U.S.	Southeast U.S.
Customer Contact	James R. Schneider, President	Steven L. Davenport, President	G. Douglas Jordan, President	Ronald L. Johnson, President
TECHNOLOGY SPECIALTIES	Comm. software; elec- tronic mail; energy mgt.; fax.; fiber optics; infra- red.; ISDN; key phone sys.; LANs; long-distance alt.; microwave; modems; muxes.; network mgt.; PBXs; see Comments	Key phone systems; LANs; long-dist. alternatives; modems; muxes; network mgt.; PBXs; T1; voice/data integration	Key phone systems; LANs; long-dist. alternatives; network mgt.; PBXs; T1; private networks; security	Comm. software; EDI; elec- tronic mail; energy mgt.; fax.; fiber optics; information security; infrared; ISDN; key phone systems; LANs; microwave; mobile radio; modems; muxes.; network management; see Comments
APPLICATIONS SPECIALTIES	Corporate; education; financial; gov't; health care; international; see Comments	Corp. users; education; financial; health care; manufacturing; multiple building environment; utilities	Corp. users; education; financial; gov't; mult. bldg. env.; real estate	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. env.; real estate; shared tenant
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	No	No	Yes
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	Yes	No	No	Yes
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Consultant	Outside companies	Consultant; outside cos.	Outside companies
Recommended Software Is Available through:	Consultant	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant	Consultant; outside cos.	Consultant; outside cos.	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude office auto.; private pay phones; radio pgng.; SNA; T1; teleconf.; telex/TWX; voice/data integ.; voice msg. Other applica- tions include manufac- turing; military; mult. bldg. env.; real estate; svc. bureaus; shared tenant svcs.; and telephone operating companies.	SCS Engineering is a pro- fessional engineering organization providing a full range of communica- tions and information management consulting services.	—	Other technologies in- clude office automation; packet switching; PBXs; PC comm.; private net- works; private pay phones; radio paging; satellite; SNA; T1; teleconf.; voice/ data integ.; voice messaging. This firm provides consulting services to many large companies in the Birmingham area.

Directory of Communications Consultants

FIRM	Spar Aerospace Ltd., Communications Group	The Steering Committee Ltd.	Systems Approach	The Tanner Group
GENERAL INFORMATION				
Address	2811 Airpark Drive Santa Maria, CA 93455	7058 Waverly Court Kansas City, KS 66109	Heritage Road RD 2 Box 314 Barnsboro, NJ 08080	60 West 200 South Suite 350 Salt Lake City, UT 84101
Telephone Number	(805) 925-2540	(913) 334-1037	(609) 222-6507	(801) 538-2320
Date Founded	1967	1988	1983	1984
Number of Active Consultants on Staff	14-20	1 plus outside cos.	Consultant did not specify	5
Geographic Area Served	International	U.S.	Philadelphia, New York City, Baltimore, Washington, DC	U.S.
Customer Contact	Thomas van der Heyden, Dir. Business Development	Mr. L. Kelly Reynolds, President	C. Yahrting, owner	Todd A. Tanner, Senior Analyst
TECHNOLOGY SPECIALTIES	ISDN; long-dist. alternatives; network mgt.; private networks; satellite; T1; teleconferencing; voice/data integration; VSAT	Communications software; electronic mail; fiber optics; info. security; ISDN; LANs; long-dist. alternatives; microwave; modems; muxes; network mgt.; office auto.; packet switch.; PBXs; PC comm.; T1; see Comments	Communications software; EDI; electronic mail; LANs; microwave; modems; muxes; network mgt.; office automation; pkt. swtchng.; PC comm.; T1	Fiber optics; infrared; ISDN; key phone systems; long-dist. alternatives; microwave; mobile radio; modems; muxes.; network mgt.; PBXs; private networks; radio paging; SNA; see Comments
APPLICATIONS SPECIALTIES	Corp. users; education; financial; gov't; health care; international; manufacturing; military; telephone cos.; utilities	Corp. users; education; financial; gov't; health care; mult. bldg. env.; utilities; data processing	Corporate users; education; health care; manufacturing; service bureaus	Corp. users; education; financial; government; international; manufact.; mult. bldg. env.; real estate; see Comments
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	No	Yes	Yes	Yes
RFP Preparation/Evaluation	No	Yes	Yes	Yes
Strategic Planning	No	Yes	No	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	No
Administrative/Management Services	No	Yes	Yes	No
Product Design and Development	No	No	No	No
Cost Analysis	No	Yes	Yes	Yes
Competitive Analysis	No	Yes	No	No
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Consultant; parent co.; outside cos.	Outside companies	Consultant; outside cos.	Outside companies
Recommended Software Is Available through:	Consultant; outside cos.	Outside companies	Consultant; outside cos.	Outside companies
Recommended Services Are Available through:	Consultant; parent co.; outside cos.	Outside companies	Consultant; outside cos.	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Principle activities center on the design of private networks requiring multiple services such as voice, data, video, etc. Principle transmission media expertise is satellite with microwave and fiber.	Services available from outside cos. include custom software; management consulting; business plans; corporate/client newsletters; annual reports; personnel; import/export transportation; and commodities transportation.	This firm specializes in graphic documentation of networks.	Other technologies include T1; voice/data integration; voice messaging; network optimization; hardware and network audits; and incoming call management. Other applications are telephone cos.; utilities; automatic call distribution; incoming call management; and transit authorities.

Directory of Communications Consultants

FIRM	Technology Concepts Inc./ A Bell Atlantic Co.	Tel-Econ Consultants, Inc.	TelCon Associates, Inc.	Telcon Associates of Chicago, Inc.
GENERAL INFORMATION				
Address	40 Tall Pine Drive Sudbury, MA 01776	26 Frazer Drive Greenlawn, NY 11740	9411 Santa Fe Drive Overland Park, KS 66212	25W041 Hobson Road, Naperville, IL 60540
Telephone Number	(508) 443-7311	(516) 261-2600	(913) 383-3200	(312) 961-0100
Date Founded	1981	1971	1973	1976
Number of Active Consultants on Staff	20	3	8	2
Geographic Area Served	U.S.	U.S.	U.S. and Canada	Chicago, northern Illinois, national
Customer Contact	Sarsha Adrian, Senior Sales Executive	Alan L. Berger, President	Bruce Thatcher, President	Richard J. Mead, President
TECHNOLOGY SPECIALTIES	Communications software; ISDN; LANs; network mgt.; private networks; and voice/data integration	Comm. software; E-mail; energy mgmt.; fax; key phone systems; LANs; long- dist. alt.; microwave; modems; muxes; network mgt.; office auto.; PBXs; PC comm.; private nets; T1; see Comments	Electronic mail; energy mgt.; fax; fiber optics; ISDN; key phone systems; LANs; long-dist. alt.; mobile radio; modems; muxes; network mgt.; PBXs; private networks; T1; see Comments	Key telephone systems; LANs; long-dist. alt.; voice/data integration; modems; muxes.; network mgt.; PBXs; private net- works; satellite; T1; telex/TWX; voice messaging
APPLICATIONS SPECIALTIES	Corp. users; financial; gov't; mult. bldg. environment and telephone companies	Corp. users; education; financial; health care; manufacturing; mult. bldg. env.; shared tenant svcs.	Corp. users; education; financial; health care; international; manufact- uring; real estate; see Comments	Corporate users; financial institutions; education; manufacturing
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	No	Yes	No	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	No	No
Administrative/Management Services	No	Yes	Yes	Yes
Product Design and Development	Yes	No	No	No
Cost Analysis	No	Yes	Yes	Yes
Competitive Analysis	No	Yes	No	No
Other	—	—	Seminar/training	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Parent co.; subsidiary; outside cos.	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Consultant; parent co.; subsidiary; outside cos.	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; parent co.; subsidiary; outside cos.	Outside companies	Consultant; outside cos.	See Comments
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	This company provides voice and data communi- cations services and software products. They specialize in connectivity, integration; and network management. Services in- clude consulting, software engineering, and product training.	Other technologies in- clude private pay phones; radio paging; satellite; teleconf.; telex/TWX; voice/data integ.; voice messaging; VSAT; and management of company telecommunications. This firm covers all aspects of the client's operation that pertains to telecom.	Other technologies in- clude private pay phones; radio paging; teleconferencing; telex/TWX; voice/data integration; voice messaging. Other applications are mult. bldg. env.; service bureaus; transportation broadcasting; whole- sale and distribution operations.	Recommended services are available from consultant, affiliated companies (TelCon Associates, Inc.), and outside companies.

Directory of Communications Consultants

FIRM	Telecom Consultants Group, Inc.	Telecom Consulting Engineers	TeleCom Management Associates	Telecom Planning
GENERAL INFORMATION				
Address	120 S. Riverside Drive Chicago, IL 60606	P. O. Box 53012 Fayetteville, NC 28305-3012	818 Jackson Street P.O. Box 2004 Columbus, IN 47202	705 Hibiscus Trail Melbourne Beach, FL 32951
Telephone Number	(312) 454-9396	(919) 864-6040	(812) 378-3133	(407) 725-9100
Date Founded	1982	1984	1984	1979
Number of Active Consultants on Staff	8	2	3	3
Geographic Area Served	U.S.	Southeast	U.S.	U.S., east of Rockies primarily
Customer Contact	Mr. William J. Mashek, President	Dan L. Shearin, owner	Not specified	John E. Dulfer, President
TECHNOLOGY SPECIALITIES	Electronic mail; fax; fiber optics; ISDN; key phone systems; LANs; long-dist. alternatives microwave; modems; muxes.; network mgt.; PBXs; private networks; T1; see Comments	Key phone systems; LANs; long-dist. alternatives; network mgt.; PBXs; private networks; and private pay phones	Office automation; key telephone systems; LANs; fiber optics; infrared; ISDN; long-distance alternatives; microwave; network mgt.; packet swtchnng.; PBXs; PC comm.; T1; see Comments	Fiber optics; LANs; modems; muxes; network mgt.; pkt. swtchnng.; PC comm.; private networks; SNA networks; T1; voice/data integration and cabling systems
APPLICATIONS SPECIALITIES	Corp. users; education; financial; government; health care; manufac- turing; multiple building environment	Corp. users; education; financial; gov't; health care; manuf.; military; mult. bldg. env.; real estate; service bureaus; shared tenant services	Corp. users; education; gov't; health care; international; manu- facturing; real estate; utils.; see Comments	Corporate users; multiple building environments; financial institutions education; government
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	No	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	No	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	No	Yes	Yes
Administrative/Management Services	Yes	Yes	Yes	No
Product Design and Development	No	No	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	No
Other	—	—	—	Computer applications design
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Outside companies	Outside companies	Consultant; outside cos.	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Other technologies in- clude private pay phones; radio paging; security; teleconferencing; video; voice/data inte- gration; voice messaging; and building and campus cable systems.	—	Other technologies in- clude private networks; private pay phones; satellite; teleconf.; voice/data integration; voice messaging; and VSAT. Other applications are shared tenant svcs.; and telephone cos. This firm specializes in consulting services for telemarketing and investigative and expert testimony for insurance companies.	Designs networks and information systems for Fortune 500 and government clients. Areas of specialization are data communications, network design, distributed computing, LANs, and integrated office communications. Former officer, STC; author; and speaker. Degrees in electrical engineering, mathematics, and economics.

Directory of Communications Consultants

FIRM	Telecom Resource Group, Ltd.	Telecom Services	Telecommunications Consulting International	Telecommunications Consulting Services
GENERAL INFORMATION				
Address	1035 East State Street Geneva, IL 60134	250 West 54th Street Suite 704 New York, NY 10019	19412 Torran Rocks Terrace Gaithersburg, MD 20879	220 Montgomery Street, #484 San Francisco, CA 94104
Telephone Number	(312) 232-6632	(212) 245-2530	(301) 330-3145	(415) 433-0150
Date Founded	1983	1984	1980	1976
Number of Active Consultants on Staff	5	Not specified	1	4
Geographic Area Served	International	International	U.S.	West Coast
Customer Contact	Patricia Lehr Director of Mktg. Svc.	S. Mazonson, President	Matthew T. Brunk	Ray Palmer Consultant
TECHNOLOGY SPECIALITIES	Comm. software; fiber optics; LANs; long-distance alt.; network mgt.; PBXs; PC comm.; private networks; private pay phones; satellite; SNA; T1; teleconf.; voice/data integration; voice messaging	Comm software; EDI; electronic mail; fax; fiber optics; info. security; infrared; interactive a/v; ISDN; LANs; long-distance alternatives; microwave; mobile radio; modems; muxes; network management; SNA; see Comments	Fax; ISDN; key phone sys.; long-dist. alt.; muxes.; network mgt.; office auto.; PBXs; private networks; T1; satellite; security; teleconf.; telex/TWX; see Comments	Communications software; electronic mail; fax.; key phone systems; LANs; long-dist. alternatives; modems; muxes.; network mgt.; PBXs; PC comm.; T1; teleconf.; telex/TWX; see Comments
APPLICATIONS SPECIALTIES	Corporate; education; financial; gov't; health care; manufacturing; real estate; mult. bldg. env.; telephone operating companies	Corp. users; education; financial; gov't; health care; international; manufacturing; military; mult. bldg. env.; service bureaus; see Comments	Corp. users; education; financial; health care; manufacturing; real est.; mult. bldg. environment; shared tenant svcs.; utilities; transportation	Corp. users; education; financial; gov't; health care; manufact.; real estate; service bureaus; shared tenant services
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	No	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	Yes	No	Yes
Administrative/Management Services	Yes	No	Yes	Yes
Product Design and Development	Yes	Yes	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	Yes	Yes
Other	—	—	Traffic engineering and troubleshooting	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Outside companies	Outside companies	Outside companies	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	—	Other technologies include office automation; pkt. swtchng.; PC comm.; private networks; private pay phones; radio paging; satellite; T1; teleconf.; video; voice/data integ.; voice messaging; VSAT; and operator services. Other applications are shared tenant svcs. and telephone companies.	T1 consult-line subscription consulting service.	Other technologies include voice/data integration and voice messaging.

Directory of Communications Consultants

FIRM	Telecommunications Engineering, Inc.	Telecommunications International, Inc.	Telecommunications Management Corporation	Telecommunications Management Inc.
GENERAL INFORMATION				
Address	1221 Abrams Suite 320, LB-20 Richardson, TX 75081	215 Union Boulevard Suite 300 Lakewood, CO 80228	200 Reservoir Street Needham Heights, MA 02194	2805 Butterfield Road Oak Brook, IL 60521
Telephone Number	(214) 644-8864	(303) 980-8993	(617) 449-1111	(312) 571-2456
Date Founded	1979	1979	1975	1973
Number of Active Consultants on Staff	5	30	11	3
Geographic Area Served	North America and international	U.S. and international	U. S.	U.S.
Customer Contact	Dr. H. Charles Baker	O. Arnold Snyder, Director of Marketing	Fred Baitl, Vice President	Larry M. Salvatori, President
TECHNOLOGY SPECIALITIES	ISDN; T1; multiplexers; local area networks; voice/data integ.; packet switching; modems; info. security; private nets.; network mgt.; communications systems integration	EDI; electronic mail; energy mgmt.; facsimile; fiber optics; ISDN; LANs; key phone systems; long-dist. alternatives; microwave; modems; muxes; network mgt.; office aut; see Comments	Fax.; fiber optics; interactive a/v; ISDN; LANs; long-dist. alt.; microwave; modems; muxes.; network mgt.; PBXs; private networks; radio paging; SNA; T1; video; voice/data integ.; voice messaging; and health care comm. sys.	Communications software; key phone systems; long- dist. alternatives; microwave; network mgt.; PBXs; private networks; voice/data integ.; voice messaging; building wiring plans; disaster planning; call processing
APPLICATIONS SPECIALTIES	Corp. users; education; financial; manufacturing; mult. bldg. env.; telephone cos.; utilities	Corp.; educ.; financial gov't; health care; int.; manufacturing corps.; multiple bldg. env.; real estate	Corp. users; education; financial; gov't; health care; mult. bldg. env.; and real estate.	Corp. users; manufacturing; mult. bldg. env.; legal firms; accounting firms; professional firms; inbound call centers
SERVICES PROVIDED				
Equipment Evaluation	No	Yes	Yes	Yes
Training/Education	Yes	No	No	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	Yes	No
Administrative/Management Services	No	Yes	Yes	Yes
Product Design and Development	No	No	Yes	No
Cost Analysis	No	Yes	Yes	Yes
Competitive Analysis	No	No	Yes	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.	Consultant	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	State-licensed engineers with average experience over 25 years in communications and computer systems.	Other technologies include pkt. swtchnng.; PBXs; PC comm.; private networks; radio paging; satellite; SNA; T1; tele- conferencing; video; voice/data integration; voice messaging; VSAT. For individ. switching, (PBX/Centrex, ACD, LAN), transmission, or dist. project, TII offers turn- key system design, pro- curement, and implemen- tation services.	TMC specializes in all areas of independent consulting. We have an in-house expertise in plant operations, MIS, and all areas of telecom. Our clients are large (1,000 lines) universities and health care facilities. Our strategic services division provides consult- ing to major manufacturers in areas of RFP response, selling through consultants.	Firm has experience with integrating call processing and automatic call distribution systems. It specializes in vendor/ customer problem solving and dispute resolution with particular expertise in the IBM/Rolm environment.

Directory of Communications Consultants

FIRM	Telecommunications Methods & Systems, Inc.	Telecommunications Network Architects	Telecommunications Systems Management, Inc.	TeleDesign Inc.
GENERAL INFORMATION				
Address	26777 Lorain Road Suite 615 North Olmsted, OH 44070	P. O. Box 1776 Safety Harbor, FL 34695	93 Centre Pointe Drive St. Charles, IL 63303	6001 Highway Blvd. Suite 1 Katy, Texas 77450
Telephone Number	(216) 734-7800	(813) 725-1444	(314) 441-6100	(713) 391-1586
Date Founded	1982	1983	Consultant did not specify	1984
Number of Active Consultants on Staff	2	5	4	10
Geographic Area Served	U.S.	International	International	Continental U.S., Canada, Pacific basin, Caribbean, S.A.
Customer Contact	David V. Johnson, President	Donald E. Kimberlin, Principal Consultant	Mike Gertken, Vice President, Marketing	John Mathis, President
TECHNOLOGY SPECIALTIES	Comm. software; EDI; electronic mail; fax.; fiber optics; infrared; interactive a/v; ISDN; key phone systems; LANs; long-dist. alternatives; microwave; mobile radio; muxes; see Comments	EDI; electronic mail; T1; energy mgt.; fax; FCC licensing; fiber optics; info. security; infrared; interactive a/v; ISDN; key phone systems; LANs; long-dist. alt.; microwave; mobile radio; modems; muxes; see Comments	Communications software; key phone systems; LANs; long-dist. alternatives; microwave; modems; network mgt.; PBXs; PC comm.; private networks; T1	Fiber optics; ISDN; key telephone systems; LANs; long-dist. alt.; network mgt.; office auto.; packet switching; PBXs; private pay phones; T1; teleconf.; voice/data integ.; voice msg.; hotel communications
APPLICATIONS SPECIALTIES	Corp. users; education; financial; gov't; health care; international; manufacturing; see Comments	Corp. users; education; financial; gov't; health care; international; manu- facturing; military; mult. bldg. env.; real estate; utils.; see Comments	Corp. users; financial; health care; manufacturing	Corporate users; health care; international; multiple building environment; hotel facilities
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	Yes	Yes	No
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	Yes	Yes	Yes
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	No	Yes	Yes	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	No	Yes	Yes	Yes
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Consultant; outside cos.	Outside companies
Recommended Services Are Available through:	Consultant	Consultant; outside cos.	Consultant; outside cos.	Consultant
Is Compensation Received from Anyone Other than the Client?	No	No	Consultant did not specify	No
COMMENTS	Other technologies include modems; network mgt.; packet switching; PBXs; PC comm.; private networks; private pay phones; radio paging; satellite; T1; teleconf.; Telex/TWX; video; voice/data integ.; voice messaging; VSAT. Other applications are mult. bldg. env.; real estate; shared tenant services.	Other technologies are network mgt.; office aut.; packet swtchnng.; PBXs; PC comm; private networks; private pay phones; radio paging; satellite; video; security; teleconf.; telex/TWX; voice/data integ.; voice messaging; VSAT; forensic cons.; disaster recov./avoid.; RFI/EMI; bldg. wiring sys. Other applications are svc. bureaus; phone cos.; broadcasters; transport.	—	—

Directory of Communications Consultants

FIRM	Telkom Associates	TelTec Services Company	D. I. Towers Consultants Ltd.	Victory and Associates
GENERAL INFORMATION				
Address	P. O. Box 1095 Bala Cynwyd, PA 19004	733 Fall Street Spring Lake, MI 49456	10 Willowood Court Willowdale, Ontario Canada M2J 2M3	311 Miller Avenue, Suite E Mill Valley, CA 94941-2844
Telephone Number	(215) 664-1884	(616) 846-1821	(416) 368-7649	(415) 383-2333
Date Founded	1984	1981	1980	1977
Number of Active Consultants on Staff	2	2	2	4
Geographic Area Served	Delaware Valley	U.S.	Canada	U.S.
Customer Contact	John Harahan	Elwyn G. Hudson, Owner	Doug Towers, President	Wallis O. Victory, President
TECHNOLOGY SPECIALITIES	LANs; modems; muxes; long-distance alt.; network management; priv. networks; SNA; T1; tele- conferencing; voice/data integration; voice messaging	Electronic mail; fax.; fiber optics; infrared; key phone systems; LANs; long-dist. alternatives; microwave; modems; muxes.; network mgt.; PBXs; private networks; T1; see Comments	Electronic mail; facsimile; key phone systems; LANs; long-dist. alternatives; mobile radio; muxes.; network mgt.; office auto- mation; PBXs; PC comm.; private networks; T1; satellite; teleconf.; see Comments	ISDN; key phone systems; long-dist. alternatives; microwave; network manage- ment; PBXs; voice/data integration; voice messaging
APPLICATIONS SPECIALTIES	Corporate users; finance; manufacturing corporations	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. env.; real estate; see Comments	Corporate users; government	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. env.; real estate; shared tenant services
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	Yes
Training/Education	Yes	No	No	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	No	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	No	Yes	No
Administrative/Management Services	Yes	Yes	Yes	Yes
Product Design and Development	Yes	No	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	No	Yes	Yes
Other	Provides consultants for above positions	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Consultant did not specify	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Consultant did not specify	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.	Consultant did not specify	Outside companies	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	This company recruits and provides consultants for voice and data users.	Other technologies in- clude private pay phones; teleconf.; voice/data integ.; voice messaging. Other applications are service bureaus; shared tenant services.	Other technologies in- clude voice/data integration and voice messaging.	—

Directory of Communications Consultants

FIRM	W and J Partnership	Wallace Communications Consultants	The Walsh Communications Group	WBC Communications
GENERAL INFORMATION				
Address	17211 Quail Court Morgan Hill, CA 95037	P.O. Box 320177 Tampa, FL 33679	2120 Spruce Street Philadelphia, PA 19103	P.O. Box 14 Jesup, MD 20794
Telephone Number	(408) 779-1714	(813) 253-2376	(215) 732-0650	(301) 498-5682
Date Founded	February 1982	1973	October 1983	1985
Number of Active Consultants on Staff	2	1	4	4
Geographic Area Served	U.S. and international	U.S.	MidAtlantic and Northeast U.S.	Maryland, Washington D.C., and Northern Virginia
Customer Contact	William A. Morgan	Irwin "Wally" Wallace	Eileen Walsh, Principal Consultant	Vincent Basignani, VP Operations; Laurence Wolfe, VP Sales
TECHNOLOGY SPECIALITIES	Fiber optics; ISDN; LANs; private networks; comm. software, info. security, long-distance alt., microwave, mobile radio, modems, muxes., network mgt. pkt. swtchnng., PBX, PC comm., see Comments	Electronic mail; fax.; fiber optics; ISDN; key phone systems; LANs; long-dist. alternatives; microwave; mobile radio; modems; muxes.; network mgt.; pkt. swtchnng.; PBXs; see Comments	Communications software; electronic mail; fax.; fiber optics; key phone systems; LANs; long-dist. alt.; microwave; network mgt.; office automation; PBXs; private networks; satellite; T1; voice/data integration; voice mssng.	ISDN; key phone systems; long-dist. alternatives; network mgt.; PBXs; priv. networks; teleconf.; video; voice/data integ.; voice messaging; call distribution systems (ACD)
APPLICATIONS SPECIALTIES	Corporate users; government agencies; international; telephone operating cos.; education; financial; See Comments	Corp. users; education; financial; health care; manufacturing; mult. bldg. env.; real estate; svc. bureaus; shared tenant	Corp. users; education; health care; manufacturing; real estate; shared tenant services	Corporate users; education; multiple building environment; shared tenant
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	Yes	No
Training/Education	Yes	Yes	Yes	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	Yes	Yes	No	No
Administrative/Management Services	Yes	No	Yes	No
Product Design and Development	Yes	No	No	No
Cost Analysis	Yes	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	No	No
Other	—	—	—	—
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside cos.	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant, outside cos.	Outside companies	Outside companies	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	Further technologies include satellite, SNA, security, private networks, voice/data integration, and VSAT. Further applications include manufacturing, military, multiple building environment, and utilities.	Other technologies include private networks; private pay phones; radio paging; T1; teleconferencing; voice/data integration; and voice messaging.	E. Walsh is a member of the Society of Telecom Consultants, which has a strong code of ethics insuring objective and non-vendor aligned consulting.	WBC Communications provides call distribution systems (ACD) expertise.

Directory of Communications Consultants

FIRM	Weaver and Associates, Inc.	Western Telecommunication Consulting Inc.	WIRM, Inc. (Wenker Information Resources Management)	W. J. P. Associates, Inc.
GENERAL INFORMATION				
Address	P.O. Box 1220 Valparaiso, IN 46384	550 S. Hill Street Suite 690 Los Angeles, CA 90013	3716 Ridgelea Drive Fairfax, VA 22031	P. O. Box 450 Rushford, NY 14777
Telephone Number	(219) 462-9581	(213) 622-4444	(703) 425-9175	(716) 437-2478
Date Founded	March 1985	1983	1981	1977
Number of Active Consultants on Staff	2	3	4	4
Geographic Area Served	Midwest	U.S.	MidAtlantic	Northeast
Customer Contact	Joann H. Weaver	Roberta Berkowitz	Dr. William J. Wenker	Richard L. Boas, Principal
TECHNOLOGY SPECIALITIES	Voice messaging; PBXs; network management	Electronic mail; fax; fiber optics; key phone systems; LANs; long-distance alternatives; microwave; network mgt.; PBXs; voice/data integration; voice messaging	Information security; ISDN; LANs; network mgt.; packet switching	Comm. software; E-mail; energy mgmt.; fiber opt; infrared; interactive a/v; ISDN; key phone sys.; LANs; long-dist. alt.; microwave; mobile radio; modems; muxes.; network mgt.; office auto.; packet switching; see Comments
APPLICATIONS SPECIALTIES	Health care; manufacturing; transportation	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. environment; shared tenant svcs.; telephone cos.	Military; educational institutions; government agencies; corporate users	Corp. users; education; financial; gov't; health care; international; manu- facturing; military; mult. bldg. env.; real estate; see Comments
SERVICES PROVIDED				
Equipment Evaluation	Yes	Yes	No	Yes
Training/Education	No	No	No	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes	Yes
Strategic Planning	No	Yes	No	Yes
Network Planning and Design	Yes	Yes	Yes	Yes
Market Research and Analysis	No	No	No	Yes
Administrative/Management Services	No	Yes	No	Yes
Product Design and Development	No	No	Yes	Yes
Cost Analysis	Yes	Yes	No	Yes
Competitive Analysis	No	Yes	Yes	Yes
Other	—	—	—	See Comments
AFFILIATIONS				
Recommended Hardware Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Software Is Available through:	Outside companies	Outside companies	Outside companies	Outside companies
Recommended Services Are Available through:	Consultant; affiliate; outside cos.;	Outside companies	Consultant	Outside companies
Is Compensation Received from Anyone Other than the Client?	No	No	No	No
COMMENTS	—	—	—	Other technologies: PBXs; PC comm.; private nets; satellite; SNA; T1; tele- conf.; video; voice/data integ.; voice messaging; VSAT; and telemarketing. Other applications include telephone companies; utilities; svc. bureaus; and shared tenant ser- vices. This company pro- vides a full range of in- formation management and movement services.

Directory of Communications Consultants

FIRM	World Communications Group	The Yankee Group	Arthur Young & Co.
GENERAL INFORMATION			
Address	One Waters Park Drive Suite 229 San Mateo, CA 94403	200 Portland Street Boston, MA 02114	515 South Flower Street Los Angeles, CA 90071
Telephone Number	(415) 570-5444	(617) 367-1000	(213) 977-3403
Date Founded	1981	1970	1896
Number of Active Consultants on Staff	3	50	20
Geographic Area Served	Primarily Western U.S.	U.S. and international	U.S. and International
Customer Contact	Peter G. Bologna, President	Howard Anderson, Managing Director Berge Ayvazian, VP	Jeanne Chapman
TECHNOLOGY SPECIALTIES	Electronic mail; fiber optics; info. security; infrared; interactive a/v; key phone systems; LANs; long-dist. alt.; microwave; mobile radio; modems; muxes.; network mgt.; see Comments	EDI; electronic mail; energy mgmt.; fax.; fiber optics; ISDN; key phone systems; LANs; long-distance alternatives; mobile radio; modems; muxes.; network mgt.; office auto.; see Comments	Electronic mail; information security; interactive audio/video; long distance alternatives; LANs; microwave; mobile radio; modems; multiplexers; See Comments
APPLICATIONS SPECIALTIES	Corp. users; education; financial; gov't; health care; manufacturing; mult. bldg. env.; real estate; svc. bureaus; utilities	Corp. users; financial; gov't; international; manufacturing; telephone cos.; utilities	Corporate users; education; financial; government; health care; utilities
SERVICES PROVIDED			
Equipment Evaluation	Yes	Yes	Yes
Training/Education	Yes	No	Yes
RFP Preparation/Evaluation	Yes	Yes	Yes
Strategic Planning	Yes	Yes	Yes
Network Planning and Design	Yes	Yes	Yes
Market Research and Analysis	No	Yes	No
Administrative/Management Services	Yes	No	Yes
Product Design and Development	No	Yes	Yes
Cost Analysis	Yes	Yes	Yes
Competitive Analysis	Yes	Yes	No
Other	—	—	Project planning; project management; system design and implementation
AFFILIATIONS			
Recommended Hardware Is Available through:	Outside companies	Consultant did not specify	Outside companies
Recommended Software Is Available through:	Outside companies	Consultant did not specify	Outside companies
Recommended Services Are Available through:	Consultant; outside cos.	Affiliate	Consultant; outside cos.
Is Compensation Received from Anyone Other than the Client?	No	No	No
COMMENTS	Other technologies include office automation; PBXs; PC comm.; private networks; private pay phones; radio paging; satellite; T1; teleconf.; telex/TWX; video; voice/data integ.; and voice messaging.	Other technologies include pkt. swtchnng.; PBXs; PC comm.; private networks; radio paging; satellite; SNA; T1; teleconf.; video; voice/data integ.; voice messaging; and VSAT.	Other technology areas include packet switching; private networks; radio paging; satellite; T1; voice/data integration; and voice messaging.

Glossary

A

A and B signaling—a procedure used in most T1 transmission facilities where one bit, robbed from each of the 24 subchannels in every sixth frame, is used for carrying dial and control information. A type of in-band signaling used in T1 transmission.

A-law encoding—encoding according to CCITT Recommendation G711, used with European 30-channel pulse code modulation (PCM) systems that comply with CCITT Recommendation G732. Employs nonuniform quantizing to obtain the desired compression characteristic.

abandoned call—call in which a caller cancels the call after a connection has been made but before conversation takes place, for example, after hearing a recorded announcement.

abbreviated and delayed ringing—allows the ringing associated with a call on a station line to be transferred to another station that has an appearance of that line. The transfer occurs automatically after two ringing cycles.

abbreviated dialing—enables a caller to dial a frequently used number by means of a few digits instead of the entire telephone number. Also called speed dialing and short-code dialing.

absorption—conversion of transmitted energy, such as an electronic signal, into heat or other forms of energy.

ACA—see *automatic circuit assurance*.

ACC—see *automatic callback calling*.

acceptance test—operating and testing of a new communications system to ensure that the system is operating satisfactorily before being accepted by the purchaser.

access—to obtain data from memory, a peripheral, or another system.

access arm—the mechanical device in a disk storage unit that holds one or more read/write heads.

access charge—cost assessed to communications users for access to the interexchange, interstate message toll telephone network to originate and receive interstate toll calls, as well as access to the customer's local access and transport area (LATA).

access code—the digit, or digits, that a user must dial to be connected to an outgoing trunk facility.

access control—hardware, software, and administrative tasks that monitor system operation, perform user identification, and record system accesses and changes.

access line—Connection from the customer to the local telephone company for access to the telephone network; also represents the connection between the serving toll center and the serving office of the interexchange carrier used for access to public switched network services. Also known as local loop.

access method—1) the technique and/or the program code in a computer operating system that provides input/output services. The access method typically carries with it an implied data and/or file structure with logically similar devices sharing access methods. Examples are IBM's VTAM and TCAM. 2) in local area networks, the technique and/or program code used to arbitrate the use of the communications medium by granting access selectively to individual stations. Examples are CSMA/CD and token passing.

access time—the interval between a request to read or store data and the completion of that task.

accounting codes (voluntary or forced)—Entered by a caller by dialing the appropriate digits for the call being placed followed by the digits composing the accounting codes. With forced accounting codes, if the caller fails to dial a valid accounting code within a specific timeframe, an intercept tone (or as a customer option, a recorded message) can be played to the caller and the call will be terminated.

accounting rate—charge per traffic unit, which can be a unit of time or information content, covering communications between zones controlled by different telecommunications authorities; used to establish international tariffs.

AC-DC ringing—telephone ringing that makes use of both AC and DC components—alternating current to operate a ringer, and direct current to aid the relay action that stops the ringing when the called telephone is answered.

ACD—see *automatic call distributor*.

ACF—see *advanced communications function*.

ACK—"acknowledge" character. A transmission control character transmitted by a station as an affirmative response to the station with which a connection has been set up. An acknowledge character may also be used as an accuracy control character. See also *NAK (negative acknowledgement character)*.

Glossary

ACK/NAK/END—acknowledged/not acknowledged/inquiry.

acoustic coupler—device that allows a telephone handset to be used for access to the switched telephone network for data transmission. Digital signals are modulated as sound waves, and data rates are typically limited to about 300 bps, some up to 1.2K bps.

ACTGA—see *attendant control of trunk group access*.

ACU—see *automatic calling unit*.

A/D—Analog/Digital.

adapter—a connective device designed to link different parts of one or more systems and/or subsystems

adaptive differential pulse code modulation (ADPCM)—encoding technique (CCITT) that allows analog voice signals to be carried on a 32K bps digital channel. Sampling is done at 8kHz with 3 or 4 bits used to describe the difference between adjacent samples.

adaptive equalization—equalization that is adjusted while signals are being transmitted in order to adapt to changing line characteristics.

adaptive routing—routing that automatically adjusts to network changes such as changes of traffic pattern or failures.

ADC—analog-to-digital conversion. A method of sampling and costing an analog signal to create a digital signal.

ADCCP—see *advanced data communications control procedure*.

add-on conference—sometimes referred to as “three-way calling.” Used in association with consultation hold features, it is a conference facility that allows the station user to add another internal party to the existing conversation.

add-on data modules—plug-in circuit boards that allow a PBX to receive and transmit both digital and analog signals.

address—1) coded representation of the destination of data, or of its originating terminal. Multiple terminals on one communications line, for example, each must have a unique address. 2) group of digits that makes up a telephone number. Also known as the called number. 3) in software, a location that can be specifically referred to in a program.

address, international telephone—code that specifies a unique address for any telephone in the world. The number excludes the international prefix but includes the country code and the national subscriber number.

ADPCM—see *adaptive differential pulse code modulation*.

advanced communications function (ACF)—a series of software products by IBM that supports sophisticated computer networking functions for IBM systems and terminals.

advanced data communications control procedure (ADCCP)—a bit-oriented protocol developed by ANSI.

advanced mobile phone service—AT&T-developed analog cellular radio standard, adopted in the USA and in many other parts of the world.

advanced private line termination (APLT)—provides the PBX user with access to all the services of an associated Enhanced Private Switched Communications Services (EPSCS) network. It also functions when associated with a Common Control Switching Arrangement (CCSA) network.

AFIPS—see *American Federation of Information Processing Societies*.

agent—computer system which provides connected terminals with a means of access to services residing in a (remote) server via RSA (remote server access).

agent sign-on/sign-off—enables any ACD agent to take ACD calls from any ACD-assigned set.

AIOD—see *automatic identification of outward dialing*.

alarm display—indicators on attendant console that show the status of the system. Usually two alarms are included: a minor alarm to signal a malfunction in a part of the system and a major alarm to indicate that system will not work at all.

alarm signals to stations—ability of a system to interface with various types of signaling systems and to transmit such alarms in the form of coded tones to all, or preselected, stations.

algorithm—a prescribed set of well-defined rules for the solution of a problem in a finite number of steps, e.g., a full statement of an arithmetic procedure for evaluating sine x to a stated precision.

all number calling (ANC)—calling by means of seven digits instead of previously used two letters plus five digits.

all trunks busy (ATB)—condition in which all trunks in a group are engaged.

allocate—to assign a resource to a specific task.

alpha test—the stage during new product research and development when a prototype is operated to ascertain that

Glossary

the system concept and design are functional and to identify areas that need further development and/or enhancement.

alphanumeric—a character set containing letters, digits, and other symbols such as punctuation marks; synonym for alphameric.

alphanumeric display for attendant console—feature in which the attendant's console provides a visual readout of the source of incoming calls. Allows the attendant to handle calls for several listed number groups easily, with the response depending on the number called.

alternate mark inversion—digital signaling method in which the signal carrying the binary value alternates between positive and negative polarities; zero and one values are represented by the signal amplitude at either polarity, while no-value "spaces" are at zero amplitude. Also called *bipolar*.

alternate routing—the routing of a call or message over a substitute route when an established route is busy or otherwise unavailable for immediate use.

alternate voice/data (AVD)—method of deriving telegraph-grade lines from unused portions of a voice circuit.

ALU—see *arithmetic logic unit*.

AM—see *amplitude modulation*.

AMA—automatic message accounting. See also *station message detail recording*.

ambient noise—communications interference present in a signal path at all times.

American Bell—name of ATTIS before AT&T was forced to drop the Bell name by the divestiture decree; see also *ATTIS*.

American Federation of Information Processing Societies (AFIPS)—an organization of computer-related societies; its members include The Association for Computer Machinery, The Institute of Electrical and Electronic Engineers Computer Group, Simulation Councils, Inc., and American Society for Information Science.

American National Standards Institute (ANSI)—a standards-forming body affiliated with the International Organization for Standardization (ISO) that develops standards for transmission codes protocols, media, and high level languages, among other things.

American Standard Code for Information Interchange (ASCII)—1) eight-bit code yielding 128 characters, both

displayed and nondisplayed (for device control); used for text transmission. 2) a widely used asynchronous protocol based on ASCII code.

amplifier—an electronic component that boosts the strength or amplitude of a transmitted, usually analog, signal; functionally equivalent to a repeater in digital transmissions.

amplitude—magnitude or size. In wave forms or signals occurring in a data transmission, a complete definition of the wave form can be made if the voltage level is known at all times. In this case, the voltage level is called the amplitude.

amplitude modulation (AM)—in communications technology, a transmission method where the height of the carrier wave is modified to satisfy the height of the signal wave; contrast with *frequency modulation*.

AMPS—see *advanced mobile phone service*.

analog—referring or pertaining to a signaling technique in which a transmission is conveyed by modulating (varying) the frequency, amplitude, or phase of a carrier. Contrast with *digital*.

analog loopback—technique for testing transmission equipment and devices that isolate faults to the analog signal receiving or transmitting circuitry, where a device, such as a modem, echoes back a received (test) signal that is then compared with the original signal. Compare with *digital loopback*.

analog signal—signal in the form of a continuous wavelike pattern representing a physical quantity, such as voltage, which reflects variations in some quantity, such as loudness of the human voice. Analog signaling is generic to the public switched telephone network (PSTN), as well as to certain other audio-frequency and radio-frequency facilities. A digital baseband signal generated by a business machine must be converted to analog form in order to transmit that signal over an analog facility, e.g., a voice-grade telephone line.

analog transmission—transmission of a continuously variable signal as opposed to a discretely variable signal. Also called analog signaling.

ANC—see *all-number calling*.

ancillary equipment—terminal or communications devices not required for the provision of basic telephone service. Answering machines, conferencing devices, and automatic dialers are types of ancillary equipment.

angle modulation—a form of modulation in which the angle of a sine wave carrier is varied.

ANI—see *automatic number identification*.

Glossary

anisochronous data channel—communications channel capable of transmitting data but not timing information. More commonly known as *asynchronous data channel*.

announcement, recorded—prerecorded spoken message played for incoming calls.

ANS—answer.

ANSI—see *American National Standards Institute*.

answer signal—supervisory signal (usually in the form of a closed loop) from the called telephone to the exchange and back to the calling telephone (usually in the form of a reverse battery) when the called number answers. Signal can also initiate call charging.

answerback—a manually or automatically initiated response from a terminal that usually includes the terminal address to verify that the correct terminal has been reached and that it is operational.

antenna—device used to transmit and/or receive radio waves. The physical design of the antenna determines the frequency range of transmission/reception.

aperture—the diameter of a parabolic antenna.

APL (a programming language)—a problem-solving programming language designed for remote terminals; it offers an extensive set of operators and data structures for handling arrays and performing mathematical functions.

APLT—see *advanced private line termination*.

append—to change or alter a file or program by adding to the end of the file or program.

application layer—the top of the seven-layer OSI model, generally regarded as offering an interface to, and largely defined by, the network user; in IBM's SNA, the end-user layer.

application program—software programs in a system are usually known as "application programs" and "supervisory programs." Application programs contain no input/output coding (except in the form of macroinstructions that transfer control to the supervisory programs) and are usually unique to one type of application.

applications technology satellite—early experimental satellites designed and launched by NASA.

ARC—see *attached resource computer*.

architecture—the physical inter-relationship between the components of a computer.

archive—a procedure for transferring information from an online storage diskette or memory area to an off-line storage medium.

ARCnet—the local networking products and philosophy developed by Datapoint Corporation.

area code—synonymous with numbering plan area (NPA); A three-digit code designating a "toll" center not in the NPA of the calling party. The first digit is any number from 2 through 9. The second digit is always a "1" or "0."

area code restriction—Ability of the switching equipment to selectively identify three-digit area codes, and either allow or deny passage of long distance calls to those specific area codes.

arithmetic logic unit (ALU)—part of a computer that performs the actual computations.

ARL—see *attendant release loop*.

ARO—after receipt of order.

ARPA—the generic name for services designed by the Department of Defense in the 1970s; it is now a de facto standard for networking multivendor computers running on differing operating systems. Includes FTP, SMTP, and TELNET.

ARQ—see *automatic repeat request*.

array—a named, ordered collection of data elements that have identical attributes; or an ordered collection of identical structures.

ARS—see *automatic route selection*.

artificial intelligence—the capability of a computer to perform such functions that are associated with human logic as reasoning, learning, and self-improvement.

artificial language—a convention based on a set of rules established prior to its usage and without a precise relationship to the user applications it will be used for. Contrast with *natural language*.

ARU—see *audio response unit*.

ASR—see *automatic send/receive*.

assembler language—a set of commands that includes symbolic machine language statements in which there is a one-to-one correspondence with computer instructions.

asynchronous—occurring without a regular or predictable time relationship to a specified event, e.g., the transmission of characters one at a time as they are keyed. Contrast with *synchronous*.

Glossary

asynchronous communications software—a software package that, when used with a modem, allows a PC to communicate with information services, databases, or hosts via asynchronous transmission over the public telephone network.

asynchronous computer—a digital system in which each operation starts as a result of signal generated by the completion of the previous operation or the availability of a device required by an operation; contrast with synchronous computer where all operations are started in step with a master clock.

asynchronous transmission—transmission in which each information character, or sometimes each word or small block, is individually synchronized, usually by the use of start and stop elements. Also called start-stop transmission.

ATB—see *all trunks busy*.

ATND—see *attendant*.

ATS—see *applications technology satellite*.

attached resource computer (ARC)—the local networking products and philosophy developed by Datapoint Corporation.

attendant (ATND)—usually refers to a local switchboard operator, for example, on a PBX. See also *Operator*.

attendant conference—feature that allows the attendant to establish a conference connection between central office trunks and internal stations.

attendant console—centralized operator position, either desktop or floor-mounted, that uses pushbutton keys for all control and call connecting functions.

attendant control of trunk group access (ACTGA)—attendant control and restriction of access by all stations to various trunk groups. Restricted calls are usually routed to the attendant.

attendant forced release—attendant-activated facility that automatically disconnects all parties on a given circuit when that circuit is entered by the attendant and the facility engaged.

attendant incoming call control—automatically diverts incoming trunk calls to a predetermined station after a pre-designated period of time (or number of rings).

attendant locked loop operation—allows the attendant at a console to retain supervision or recall capability of any particular call that was processed.

attendant lockout—denies an attendant the ability to re-enter an established trunk/station connection unless recalled by the station.

attendant loop transfer—allows the attendant to transfer any call to another attendant for processing.

attendant monitor—special attendant circuit that allows listening in on all circuits with the console handset/headset transmitter deactivated.

attendant override—allows an attendant to enter a busy trunk connection and key the trunk number within the PBX. A warning tone will be heard by the connected parties, after which the connected parties and the attendant will be in a three-way connection.

attendant recall—calls held, camped-on, or completed by the attendant are returned to the console if unanswered within a preset period.

attendant release loop—incoming trunk and station calls extended by the attendant to idle stations can be released from the console switched loop when the attendant presses the release key. This feature is used for releasing incoming calls to called stations that are in a busy or unanswered ringing state.

attendant repertory dialing—single pushbutton selection of repeatedly used preprogrammed numbers for outgoing calls made by the attendant on central office trunks.

attendant restriction—attendant denied the ability to gain access to a trunk in order to originate a call (or to prevent "listening in") unless assistance is required by station recall.

attendant station busy lamp—when the desired station number is keyed by the attendant, a lamp lights, providing a positive visual busy indication.

attendant station number display—unit on the attendant's console that displays the station calling number on an attendant trunk.

attendant supervisory console—special attendant console used by the chief operator of large PBX systems. It provides, in addition to standard console operation, certain monitoring and control functions over other consoles.

attendant transfer of incoming call—allows the station user connected to an incoming outside call to signal the attendant by flashing the switchhook to request the attendant to transfer the call to another station.

attendant transfer, outgoing—allows the attendant to be recalled and to transfer an outgoing exchange call originated by an internal station.

Glossary

attendant trunk busy lamp field—special panel to indicate the busy condition of each trunk circuit, one lamp being associated with each circuit.

attendant trunk group busy lamps—provides the attendant at a console with a visual display when all trunks in a given trunk group are busy.

attended operation—situation in which both data stations require attendants to establish the connection and transfer the data sets from talk (voice) mode to data mode.

attenuation—a decrease in the power of a received signal due to loss through lines, equipment, or other transmission devices. Usually measured in decibels.

ATTIS—AT&T Information Systems; division of AT&T Technologies that supplied and manufactured customer premises equipment ca. 1984-86.

audible ringing tone—tone received by the calling telephone indicating that the called telephone is being rung (formerly called ringback tone).

audio frequencies—frequencies that correspond to those that can be heard by the human ear (usually 30 Hz to 20,000 Hz).

audio response unit (ARU)—device that provides prerecorded spoken responses to digital inquiries from a telephone caller once the connection is established. Requires use of a pushbutton telephone; callers are instructed to press a certain number to obtain a certain type of information.

audiotex—communications system that allows a host computer to pass data to a voice mail computer, a store-and-forward mechanism for digitized voice, where it is interpreted and delivered over the telephone.

authorization code—code allowing a station user to override the restriction level associated with that user's station line or incoming trunk.

auto answer—automatic answering; capability of a terminal, modem, computer, or similar device to respond to an incoming call on a dial-up telephone line and to establish a data connection with a remote device without operator intervention. Allows unattended operation for incoming dial-up calls.

auto call—automatic calling; a machine feature that allows a transmission control unit or a station to automatically initiate access to a remote system over a switched line.

auto dial—automatic dialing; the capability of a terminal, modem, computer, or a similar device to place a call over the switched telephone network and establish a connection without operator intervention.

autoexecute—an operating system utility that allows the program to automatically run a program.

automated attendant system—processor controlled system that performs most attendant functions, such as answering calls, extending them to station users, taking messages, and providing assistance. The system provides voice prompts, to which users reply via any standard 12-button dialpad.

automatic call distributor (ACD)—system designed to distribute a large volume of incoming calls uniformly to a number of operators or agents, e.g., for airline reservations.

automatic callback calling (ACC)—feature that allows the station user, when encountering an internal station busy signal, to dial a 1- or 2-digit code and hang up. When both parties are free, the system automatically rings and connects the parties. While activated, this feature does not prevent the calling station from either initiating or receiving calls.

automatic calling unit (ACU)—a unit, which may or may not be integrated within a modem, that automatically dials calls based on digits supplied by the attached business machine.

automatic circuit assurance (ACA)—assists the customer in identifying possible trunk malfunctions. The PBX maintains a record of the performance of individual trunks relative to short hold time (SHT) and long hold time (LHT) calls. A significant increase in the number of short calls or a single long call can indicate a trunk failure. When a possible failure is detected, a referral call is initiated to the attendant.

automatic dialer—device that allows the user to dial pre-programmed numbers by pushing a single button. Also called an auto dialer.

automatic dialing—feature button on specialized key sets that can be assigned to dial a user-designated telephone number when pressed.

automatic dialing unit—device capable of automatically generating dialing digits.

automatic error correction—transmission system feature whereby a proportion of errors in the received signal are detected and automatically corrected.

automatic exclusion—first station accessing a line automatically prevents any other station from gaining access to that line.

automatic hold—allows the attendant to go from one trunk call to another without using a "hold" button.

automatic identification of outward dialing (AIOD)—hardware system or PBX feature that automatically obtains

Glossary

the identity of a calling station over a separate data link for the purpose of automatic message accounting.

automatic interflow—incoming ACD call reroutes from the primary gate to which it is assigned to a PBX hunt group or an alternative ACD system when the queue time exceeds a particular time interval.

automatic intraflow—automatic rerouting of an incoming ACD call from the primary gate to which it is assigned to an alternative split within the same ACD.

automatic line hold—as long as a station does not go on-hook, activation of various line pushbuttons automatically places the first line in a hold condition without the use of a special "hold" button.

automatic message accounting (AMA)—automatic recording system that documents all the necessary billing data of subscriber-dialed long distance calls. See also *station message detail recording*.

automatic message-switching center—in a communications network, location at which data is automatically routed according to its destination.

automatic number identification (ANI)—feature that automatically identifies a calling station; for use in message accounting.

automatic recall—1) automatic alert of the attendant, after a prescribed period of time, to a camped-on or unanswered call so the attendant can provide a status report to the calling outside party. 2) ability of a terminal to automatically attempt to call back a busy terminal in order to establish a call when the called terminal is free.

automatic redial—allows any station or trunk circuit to retry a connection to a busy station automatically, or, after a given number of attempts to that busy station, to be rerouted to an intercept or alternate station if the originally desired station does not become free within a prescribed period of time.

automatic repeat request (ARQ)—an error control technique that requires retransmission of a data block that contains detected errors. A special form, called "go-back-n," allows multiple blocks to be acknowledged with a single response. "Stop and wait" requires an acknowledgment after each block.

automatic route selection (ARS)—provides automatic routing of outgoing calls over alternative customer facilities based on the dialed long-distance number. The station user dials either a network access code or a special ARS access code followed by the number. The PBX routes the call over the first available special trunk facility, checking in a customer-specified order. Alternative routes can also include tie trunks to a distant PBX. When such routing is

used, the restriction level associated with the call can be transmitted to the distant PBX as a traveling class mark. Incoming tie trunks from other locations (e.g., main or satellite) can be arranged to have automatic access to ARS. This allows station users at these systems to dial a single access code to use the ARS feature at a distant PBX.

automatic send/receive (ASR)—a teleprinter terminal with off-line storage capabilities (e.g., a paper tape or magnetic recording device) that permits a message to be originated off-line for later transmission. Received messages can be accepted at speeds faster than the printer operates and stored for later printing off-line. May be used on-line as well. Also called *buffered automatic send/receive (BASR)*. Contrast with *keyboard send/receive*.

automatic station restriction—prevents unauthorized (and unaccountable) phone calls from vacant hotel rooms by automatically activating Controlled Outward Restriction when the guest is checked out from the room and deactivating the restriction when a guest is checked in.

automatic time-out on uncompleted call—switching equipment will automatically connect a station to an intercept position if the station stays off-hook without dialing for a predetermined time interval or stays connected to a busy signal longer than the predetermined time interval.

AUTOVON—automatic voice network (U.S. military).

Aviane—a heavy launch vehicle produced by the European Space Agency (ESA).

AVD—see *alternate voice/data*.

AVD circuits—alternate voice/data circuits that have been conditioned to handle both voice and data traffic.

B

babyphone—feature allowing calls to an off-hook telephone to listen to room noises, for example, to check if a baby is crying.

backbone network—the portion of a communications facility that connects primary nodes; a primary shared communications path that serves multiple users via multiplexing at designated jumping-off points.

background music—optional facility provided through the switching equipment and associated with special stations with paging speakers.

backhaul—the terrestrial link between an earth station and a switching or data center.

backup—the provision of facilities, logical or physical, to speed the process of restart and recovery following failure.

Glossary

balanced-to-ground—with a two-wire circuit, the impedance-to-ground on one wire equals the impedance-to-ground on the other wire. Compare with *unbalanced-to-ground*, a preferable condition for data transmission.

balun—balanced/unbalanced. In the IBM Cabling System, refers to an impedance-matching device used to connect balanced twisted-pair cabling with unbalanced coaxial cable.

band—1) the range of frequencies between two defined limits. 2) in relation to WATS service, the specific geographical area which the customer is entitled to call.

bandpass filter—circuit designed to allow a single band of frequencies to pass, neither of the cut-off frequencies being zero or infinity.

bandwidth—the range of frequencies, expressed in hertz (Hz), that can pass over a given transmission channel. The bandwidth determines the rate at which information can be transmitted through the circuit. The greater the bandwidth, the more information that can be sent through the circuit in a given amount of time.

baseband—pertaining or referring to a signal in its original form and not changed by modulation. A baseband signal can be analog (e.g., originating from a telephone set) or digital (e.g., originating from a business machine).

baseband or basic signal—original signal from which a transmission waveform can be produced by modulation. In telephony it is the speech waveform.

baseband signaling—transmission of a digital or analog signal at its original frequencies, i.e., a signal in its original form, not changed by modulation. It can be an analog or digital signal.

base station—within a mobile radio system, a fixed radio station providing communication with mobile stations and, where applicable, with other base stations and the public telephone network.

Basic—beginner's all-purpose symbolic instruction code. A simplified language used in programming computers.

basic rate interface—in ISDN, the interface to the basic rate CCITT 2B + D, 2 channels + 1 signaling channel. See *integrated services digital network*.

basic telecommunications access method (BTAM)—IBM's lowest level I/O macro-routine support for providing communications programs on a host computer.

BASR—buffered automatic send/receive; see *automatic send/receive*.

batch mode—application programs run on the computer one at a time. For example, financial transactions may be accumulated for a week, then fed as a group into the computer to update the general ledger files and to produce accounting reports.

batch processing—a technique in which a number of similar data or transactions are collected over a period of time and aggregated (bunched) for sequential processing as a group during a machine run.

battery reserve power—provides an alternative, independent source of power to maintain PBX service during a power failure or "brownout" at the customer location. The power supply consists of storage batteries and permanently connected battery chargers operating from a commercial AC power source.

baud—a measure of data rate, often used to denote bits per second (bps). A baud is equal to the number of discrete conditions or signal events per second. There is disagreement over the appropriate use of this word since, at speeds above 2400 bps, the baud rate does not equal the data rate in bits per second.

baudot code—a data code using a five-bit structure used on vintage teleprinter (e.g., Telex) terminals. Two of the possible 32-bit patterns are used to indicate case shift, which doubles the interpretations of the remaining bit patterns. An additional four characters are reserved for control functions, limiting the number of data characters to 52.

bay—see *rack*.

BCC—see *block check character*.

BCD—see *binary coded decimal*.

BDC—see *block down conversion*.

BDLC—see *Burroughs Data Link Control*.

beam width—angular width of an antenna radiation pattern, or beam, within which the radiation exceeds some proportion of the maximum.

Bell Operating Company (BOC)—any of 22 local telephone companies spun off from AT&T as a result of divestiture, such as Bell of Pennsylvania, New Jersey Bell, and Southern Bell. Does not include Southern New England Telephone or Cincinnati Bell. The 22 operating companies are divided into seven regions and are held by seven RBHCs (Regional Bell Holding Company).

Bellcore—Bell Communications Research; organization established by the AT&T divestiture, representing and funded by the RBHCs, for the purpose of establishing telephone-network standards and interfaces; includes much of former Bell Labs.

Glossary

benchmark—a point of reference from which measurements can be made; involves the use of typical problems for comparing performance and is often used in determining which computer can best serve a particular application.

BER—see *bit error rate*.

Berkeley (Berkeley Services)—a group of services designed by U.C. Berkeley for the educational community, networked running UNIX. Includes *rcp*, *rlogin*, *rmsh*, *rwho*, *ruptime*, and *sendmail*.

Berkeley Software Distribution Sockets (BSD)—corresponds to the session layer (5) of the OSI model. Inter-process communications for creating distributed application programs.

BERT—see *bit error rate test/tester*.

beta test—the stage at which a new product is tested under actual usage conditions.

BH—see *busy hour*.

bibliographic database—this type of reference database contains citations to published literature. Abstracts are often included with these citations to newspapers, journals, books, monographs, patents, conference proceedings, contracts, radio or television transcripts, and dissertations. See *database* and *reference database*.

bidirectional printing—printing output in two directions—left to right and right to left. This is faster and saves wear on the printer.

Big Blue—a nickname for IBM. Originally all IBM computers came in blue cabinets, hence the nickname.

binary—the base-2 number system using only the symbols 0 and 1. Since 0 and 1 can be represented as on and off, or negative and positive charges, most computers do their calculations in binary.

binary code—representation of quantities expressed in the base-2 number system.

binary coded decimal (BCD)—a binary-coded alphanumeric notation in which each of the decimal digits is represented by a binary numeral, e.g., in binary coded decimal notation that uses the weights 8-4-2-1, the number 23 is represented by 0010 0011 (compare its representation 10111 in the pure binary numeration system).

binary phase-shift keying—see *phase shift keying*.

binary synchronous communications (BSC)—bisync; a half-duplex, character-oriented data communications protocol originated by IBM in 1964. It includes control characters and procedures for controlling the establishment of a

valid connection and the transfer of data. Although still enjoying widespread usage, it is being replaced by IBM's more efficient protocol, Synchronous Data Link Control (SDLC).

bipolar—1) the predominant signaling method used for digital transmission services, such as DDS and T1, in which the signal carrying the binary value successfully alternates between positive and negative polarities. Zero and one values are represented by the signal amplitude at either polarity, while no-value "spaces" are at zero amplitude. 2) a type of integrated circuit that uses both positively and negatively charged currents, characterized by high operational speed and cost. Also called *alternate mark inversion*.

bipolar violation—modified bipolar signaling in which a control code is inserted into the original digital format.

BISDN—broadband ISDN. See *integrated services digital network*.

bisync—see *binary synchronous communications*.

bit—contraction of "binary digit," the smallest unit of information in a binary system. A bit represents the choice between a mark or space (one or zero) condition.

bit duration—equivalent to the time that it takes one encoded bit to pass a point on the transmission medium. In serial communications, a relative unit of time measurement used for comparison of delay times (e.g., propagation delay, access latency), where the data rate of a transmission channel can vary.

bit error rate (BER)—in data communications testing, the ratio between the total number of bits transmitted in a given message and the number of bits in that message received in error. A measure of the quality of a data transmission, usually expressed as a number referred to a power of 10; e.g., 1 in 10^5 .

bit error rate test/tester (BERT)—a test conducted by transmitting a known pattern of bits (commonly 63, 511, or 2047 bits in length), comparing the pattern received with the pattern transmitted, and counting the number of bits received in error. Also see *bit error rate*. Contrast with *block error rate test/tester (BLERT)*.

bit-oriented—describes a communications protocol or transmission procedure where control information is encoded in fields of one or more bits. Oriented toward full-duplex link operation.

bit-mapped—refers to a display screen on which a character or image is generated and refreshed according to a binary matrix (bit map) at a specific location in memory.

bit rate—the speed at which bits are transmitted, usually expressed in bits per second (bps).

Glossary

bits per second (bps)—basic unit of measurement for serial data transmission capacity; abbreviated as K bps, or kilobit/s, for thousands of bits per second; M bps, or megabit/s, for millions of bits per second; G bps, or gigabit/s for billions of bits per second; T bps, or terabit/s, for trillions of bits per second.

blanking interval—the area in a video signal that falls between frames. It is often used to accommodate data including synchronizing information.

BLERT—see *block error rate test/tester*.

BLF—see *busy lamp field*.

block—a string of records, words, or characters treated as a logical entity. Blocks are separated by interblock gaps, and each block may contain one or more records.

block check character (BCC)—a control character appended to blocks in character-oriented protocols. Used for determining if the block was received in error in longitudinal and cyclic redundancy checking.

block down conversion—the conversion of a full satellite band to a lower frequency, (e.g., from SHF to UHF or VHF).

block error rate—in data communications testing, the ratio between the total number of blocks transmitted in a given message and the number of blocks in that message received in error; a measure of the quality of a data transmission.

block error rate test/tester (BLERT)—a test conducted by transmitting a known blocked bit pattern, comparing the pattern received with the pattern transmitted, and counting the number of blocks containing errored bits. Also see *block error rate*. Contrast with *bit error rate test/tester (BERT)*.

block length—a measure of the size of a block, usually specified in units such as records, words, computer words, or characters.

BNA—see *Burroughs Network Architecture*.

board—the circuit card on which integrated circuits are mounted.

BOC—Bell Operating Company.

bootstrap loader—an input routine in which simple preset computer operations are used to load instructions that in turn cause further instructions to be loaded until the complete computer program is in storage. The term refers to the system “pulling itself up by its bootstraps.”

bottleneck—the operation with the least capacity in a total system with no alternative routings; the total system can be effectively scheduled by simply scheduling the limiting operation.

bps—see *bits per second*.

break—an interruption to a transmission; frequently a provision permitting a controlled terminal to interrupt the controlling computer.

breakout box—a device that allows access to individual points on a physical interface connector (e.g., EIA RS-232-C) for testing and monitoring.

breakout panel—a breakout box mounted as a component in some larger device.

bridge—1) to connect a load across a circuit. 2) the connection itself that allows the interconnection of LANs, permitting communication between devices on separate networks.

bridged ringing—system in which ringers on a line are connected across that line.

bridge lifter—device that removes, either electrically or physically, bridged telephone pairs. Relays, saturable inductors, and semiconductors are used as bridge lifters.

broadband—1) transmission equipment and media that can support a wide range of electromagnetic frequencies. 2) any voice communications channel having a bandwidth greater than a voice grade telecommunications channel; sometimes used synonymously with wideband. 3) typically the technology of CATV (q.v.) transmission, as applied to data communications; employs coaxial cable as the transmission medium and radio frequency carrier signals in the 50M Hz to 500M Hz range.

broadcast—a transmission to multiple receiving locations simultaneously. A broadcast can be made, for example, over a multipoint line to all terminals that share the line, or over a radio or television channel to all receivers tuned to that channel.

broadcasting service—radio communications service of transmission to be received directly by the general public. It can consist of sound, video, facsimile, or other one-way transmission.

brownout operation—in response to heavy demand, main system voltages are sometimes lowered leading to brownouts, where power is not lost but reduced. Although conventional PBX equipment is relatively immune to brownouts, the computer controlling the system is very sensitive to voltage variations and could fail under these conditions. Most PBXs today have the capability to cope with these reductions, or a heavy-duty power supply can be

Glossary

furnished as an option. An uninterruptible power supply (UPS) can be installed to ensure continued service during prolonged outages.

BSC—see *binary synchronous communications*.

BT—see *busy signal*.

BTAM—see *basic telecommunications access method*.

buffer—storage device used to compensate for a difference in rate of data flow, or time of occurrence of events, when transmitting data from one device to another.

bug—a mistake or malfunction.

bundled—a pricing strategy in which a computer manufacturer includes all products—hardware, software, services, training, etc.—in a single price.

Burroughs Data Link Control (BDLC)—a bit-oriented protocol similar to high-level data link control (HDLC). (Burroughs is now Unisys.)

Burroughs Network Architecture (BNA)—a communications architecture developed by Burroughs. (Burroughs is now Unisys.)

burst—in data communications, a sequence of signals counted as one unit in accordance with some specific criterion or measure.

bus—1) physical transmission path or channel. Typically an electrical connection, with one or more conductors, wherein all attached devices receive all transmissions at the same time. 2) local network topology, such as used in Ethernet and the token bus, where all network nodes listen to all transmissions, selecting certain ones based on address identification. Involves some type of contention-control mechanism for accessing the bus transmission medium.

business television—a form of business communications that enables organizations to transmit corporate television programs to reach audiences in many dispersed locations.

busy (BY)—describes a line or trunk that is in use.

busy hour (BH)—continuous one-hour period that has the maximum average traffic intensity.

busy lamp—visual indicator on a piece of telephone equipment that indicates the associated line is in use.

busy lamp field (BLF)—panel providing the attendant with visual indications of either busy or idle conditions for a particular group of one hundred station lines selected by the attendant.

busy override or intrude—see *busy verification of station lines; executive busy override*.

busy signal—1) audible and/or flashing signal that indicates that the called number is unavailable. Also called engaged tone. 2) signal that indicates all voice paths are temporarily unavailable.

busy tone (BT)—see *busy signal*.

busy verification of station lines—attendant confirmation that a line is actually in use by establishing a connection to an apparently busy line. Prior to the attendant's connection the PBX sends a burst of tone to the line to alert the talking parties.

butn—*button*.

BY—*busy*.

bypass—the name given to the method of establishing a communication link without using the facilities of the local exchange carrier (telephone company).

byte—small group of bits of data that is handled as a unit.

C

C—a highly structured portable language.

C band—portion of the electromagnetic spectrum, approximately 4G Hz to 6G Hz, used primarily for satellite and microwave transmission.

C conditioning—type of line conditioning that controls attenuation, distortion, and delay distortion so they lie within specific limits.

cable—assembly of one or more conductors within a protective sheath; constructed to allow the use of conductors separately or in groups.

CACS—see *customer administration center system*.

CAD/CAM—Computer-Aided Design/Computer-Aided Manufacturing; systems that aid in the design of products and then transfer the information to control manufacturing equipment.

CALC—customer access line charges (Centrex).

call—any demand to set up a connection. Also used as a unit of telephone traffic.

call accounting system—device that tracks outgoing calls and records data for reporting. See *station message detail recording*.

call detail recording (CDR)—see *station message detail recording*.

Glossary

call diversion—automatic switching of a call from the number to which it was directed to another predetermined number.

call duration—interval of time between the establishment of the connection between the calling and called stations and the termination of the call.

call forwarding (CF)—service or feature that allows a call to be forwarded to a number other than the one dialed.

call forwarding, all calls (CAFC)—user instruction to the system to forward calls directed to his or her station to another number. The feature is usually activated by dialing an access code followed by the number to which calls are to be forwarded. In certain systems, calls can only be directed to other stations within that system, although some PBXs allow calls to be forwarded to points outside the system.

call forwarding, busy line—automatic rerouting of incoming calls directly to the attendant or predetermined secondary station when the called station is busy.

call forwarding, don't answer—automatic rerouting to a secondary station or attendant when a given station does not answer within a prescribed time interval. (The exact interval depends on the type of switching system but is generally after three rings.)

call hold—allows a station user to hold any call in progress on that station line by flashing the switchhook, dialing a second digit such as "1", or depressing a special station push-button that will automatically provide a second dial tone for the purpose of originating another call.

call holding time (CHT)—total length of time that a circuit is held busy.

call information logging—automatic recording of chargeable calls made on a PBX system, with details of extension number, exchange line number, time, call duration, and digits dialed. This can be used for call accounting or billing.

call pickup—ability of a station user to answer any call directed to another station within a given pickup group by dialing a pickup code from either an idle or a busy state. If more than one call is waiting to be picked up in a group, the call to be answered is selected randomly. The user of an electronic telephone can activate Call Pickup by depressing the assigned button when a station line within the same pickup group is ringing. When a line in the pickup group is ringing, the Call Pickup status lamp will flash. If activated while busy, the present call will automatically be placed on Call Hold.

call processing—sequence of operations performed by a switching system from the acceptance of an incoming call through the final disposition of the call.

call progress tones—audible signals returned to the station user by the switching equipment to indicate the status of a call; dial tones and busy signals are common examples.

call record—all recorded data pertaining to a single call.

call redirection—allows redirection of calls between terminals.

call restriction—PBX feature that prevents selected extension stations from dialing external calls or reaching the operator except through the PBX attendant.

call splitting—feature that allows an attendant to speak privately to either party in a connection without the other party hearing.

call transfer—feature that allows the calling or called customer to instruct the switching equipment or operator to transfer a call or calls to another station.

call waiting (CW)—indication by a lamp on the attendant's position that incoming calls are in a queue. A steady lamp indicates that less than a specified number of calls (usually five) are waiting, and a flashing lamp indicates that more than five calls are waiting.

call waiting service—call directed to a busy station is held while a special tone is directed to the busy station user. The station user can then answer this waiting call by hanging up and then being signaled by the waiting call. Alternatively, the user can depress the switchhook and put the first call on hold by dialing the hold code, and then answer the waiting call.

called party—the subscriber requested by the calling subscriber. Also known as *called subscriber*.

called subscriber—see *called party*.

calling number display to attendant—provides the attendant with the number of the inside station that has been connected either by dialing "0" or through interception. Some systems also display the station's class of service.

calling number display to station—provides a called station with a display of the number of the calling station within the PBX. This feature generally requires additional equipment for implementation.

calling party—the subscriber who originates a call. Also known as *calling subscriber*.

calling relay—relay that is controlled via a subscriber's line or trunk line. Also called a *line relay*.

calling subscriber—see *calling party*.

CAMA—see *centralized automatic message accounting*.

Glossary

camp-on—feature whereby a subscriber calling a busy number is placed in a waiting condition; both phones ring automatically when the called party hangs up.

CAP—see *customer administration panel*.

card (circuit)—the individual boards that carry the necessary circuits for particular functions; these cards (or boards) are designed to fit expansion slots provided by many computer manufacturers.

card cage—a frame for holding circuit cards in a microprocessor; a standard cage holds 9 cards; units with motherboards can hold up to 20 cards; also referred to as a card chassis.

carriage return (CR)—a control character that causes the print or display position to move to the first position on the next line.

carrier (CXR)—a signal of known characteristics (for example, frequency) that is altered (modulated) to transmit information. Knowing the expected signal, the receiving terminal interprets any change in signal as information. Changes to the signal made by outside influences (noise) can cause the receiving terminal to misinterpret the information transmitted. See also *common carrier* and *specialized carrier*.

carrier frequency—frequency of the carrier wave that is modulated to transmit signals.

carrier sense multiple access (CSMA)—a local area network access technique in which multiple stations connected to the same channel are able to sense transmission activity on that channel and to defer the initiation of transmission while the channel is active. Similar to contention access.

carrier sense multiple access with collision detection (CSMA/CD)—a refinement of CSMA in which stations are able to detect the interference caused by simultaneous transmissions by two or more stations (collisions) and to retransmit colliding messages in an orderly manner.

carrier signaling—any of the signaling techniques used in multichannel carrier transmission. The most commonly used techniques are in-band signaling, out-of-band signaling, and separate channel signaling.

carrier system—means of obtaining a number of channels over a single path by modulating each channel on a different carrier frequency and demodulating at the receiving point to restore the signals to their original frequency.

Carterfone decision—landmark 1968 FCC decision that first permitted the electrical connection of customer-owned terminal equipment to the telephone network.

CAS—see *centralized attendant service*.

cathode-ray tube (CRT)—a device similar to a television picture tube, used to display textual and graphic information. Also called a *video display terminal (VDT)*.

CATV—see *community antenna television*.

CBX—1) computerized branch exchange. 2) centralized branch exchange. See *private automatic branch exchange*.

CCIR—Comité Consultatif International de Radiocommunication. Technical committee set up under the international telecommunications union (ITU) with responsibility for radio communications.

CCIS—see *common channel interoffice signaling*.

CCITT—Comité Consultatif International de Téléphonie et de Télégraphie. An advisory committee to the international telecommunications union (ITU) whose recommendations covering telephony and telegraphy have international influence among telecommunications engineers, manufacturers, and administrators.

CCS (hundred call seconds)—unit of telephone traffic load calculated by multiplying the number of calls per hour by the average call duration in seconds and dividing the result by one hundred (e.g., 10 CCS = 1,000 seconds).

CCSA—see *common-control switching arrangement*.

CCSA access—provision of inward and outward service between the PBX and the CCSA (common-control switching arrangement) network.

CCSR—see *centrex customer station rearrangement*.

CCTV—see *closed circuit television*.

cellular radio—technology employing low-power radio transmission as an alternative to local loops for accessing the switched telephone network. Differs from standard mobile telephony in that service is provided through a large number of areas or cells which are served by a low-power transmitter in each cell, rather than through a single high-power transmitter for the entire region. Because any given frequency can be re-used in each cell, the number of subscribers who can be served is multiplied dramatically. Users can be stationary or mobile, in which case they are passed under control of a central site from one cell's transmitter to an adjoining one with minimal switchover delay.

CEN—European Standards Institute (Comite Europeen de Normalisation).

CENLEC—European Electrical Standards Institute (Comite Europeen de Normalisation Electrique).

Glossary

central office (CO)—the physical location where communications common carriers terminate customer lines and locate the switching equipment which interconnects those lines. See also *exchange*.

central processing unit (CPU)—computer circuitry controlling the interpretation and execution of instructions.

centralized attendant service (CAS)—a PBX feature that allows a group of attendants at one location to answer and service calls for a number of locations. This arrangement works on a *switched release loop* principle, where the call, arriving on a trunk at location A, is extended through a tie line to the attendant at location B. The attendant answers the call and, using a tie line acting as a data link between the console and the PBX at location B, instructs the switch at that location to complete the call. When the attendant releases the call, the tie line between A and B is freed to handle another call.

centralized automatic message accounting (CAMA)—automatic message accounting system that is located at an exchange but that serves various adjacent exchanges. Calls not processed by ANI (Automatic Number Identification) must be routed through an operator who dials the calling number into the equipment.

Centrex—central exchange; the telephone company's termination point for multiple lines from customers. Switching facilities at the central exchange allow interconnection of the various lines or circuits.

Centrex customer station rearrangement (CCSR)—feature that allows Centrex users to make their own moves and changes. This feature requires the use of customer premises equipment (CPE).

CF—see *call forwarding*.

CFAC—see *call forwarding, all calls*.

chad—material removed when forming a hole or notch in a storage medium such as punched tape or punched cards.

chadless tape—perforated tape with the chad partially attached like a hinged flap to facilitate interpretive printing on the tape.

channel—1) in data communications, a one-way path along which signals can be sent between two or more points. Contrast with *circuit*. 2) in telecommunications, a transmission path (may be one-way or two-way, depending on the channel) between two or more points provided by a common carrier. See also *link, line, circuit, or facility*.

channel bank—equipment typically used in a telephone central office that performs multiplexing of lower speed, digital channels into a higher speed composite channel. The channel bank also detects and transmits signaling informa-

tion for each channel and transmits framing information so that time slots allocated to each channel can be identified by the receiver.

channel capacity—an expression of the maximum data traffic that can be handled by the channel.

channel service unit (CSU)—a component of customer premises equipment (CPE) used to terminate a digital circuit, such as a DDS or T1 facility, at the customer site. Performs certain line-conditioning functions, ensures network compliance per FCC rules, and responds to loopback commands from the central office. Also, ensures proper 1's density in a transmitted bitstream and performs bipolar violation correction.

channel, voice grade—channel suitable for transmission of speech, digital or analog data, or facsimile, generally with a frequency range of about 300 Hz to 3000 Hz.

character—letter, figure, number, punctuation, or other sign contained in a message. Because data is handled and transferred as a series of characters, the term is also used to mean one bit pattern in a specific data code.

character generator—the subsystem in a display unit or printer that creates symbols from the codes used to represent them.

character-oriented—describes a communications protocol or transmission procedure that carries control information encoded in fields of one or more bytes.

check-in/check-out—feature allowing the hotel console operator to activate or deactivate all hotel service related to a guest room. When check-in service is activated, the guest room station is enabled for unrestricted use and all room status information is cleared. When check-out service is activated, the guest room telephone is restricted from originating outgoing calls (room cut-off) and the printer (option) will automatically print out room status and message registration information.

chip—the substrate upon which VLSI/LSI circuits are fabricated; sometimes used to refer to the circuits on the chip themselves.

chip set—a group of microprocessors required to expand computer memory to a given increment specified by the manufacturer; e.g., for most PCs, eight 64-bit chips comprise a set offering 64K bytes of RAM.

CHT—see *call holding time*.

CICS—see *customer information control system*.

Glossary

circuit—1) means of two-way communication between two or more points. 2) a group of electrical/electronic components connected to perform a specific function. See also *channel*.

circuit board—a flat card with connections for integrated circuits.

circuit, four-wire—communications path in which four wires (two for each direction of transmission) are connected to the station equipment.

circuit grade—the data-carrying capability of a circuit; the grades of circuit are broadband, voice, subvoice, and telegraph.

circuit switching—temporary direct connection of two or more channels between two or more points in order to provide the user with exclusive use of an open channel with which to exchange information. A discrete circuit path is set up between the incoming and outgoing lines, in contrast to message switching and packet switching, in which no such physical path is established. Also called *line switching*.

circuit terminating arrangements (CTA)—arrangement for terminating activated circuits at customers' premises. May involve changes to presentation, e.g., 4-wire to 2-wire on wideband circuits.

circuit, two-wire—circuit formed by two conductors insulated from each other that can be used as a one-way transmission path, a half-duplex path, or a duplex path.

ckt—circuit.

cladding—in fiber optic cable, a colored, low refractive index material that surrounds the core and provides optical insulation and protection to the core.

class of exchange—ranking assigned to switching point in the telephone network determined by its switching functions, interrelationships with other exchanges, and transmission requirements. Also called *class of office*.

class of service—1) in X.25 networks, refers to the speed at which the data circuit terminating equipment (DCE) and the data terminal equipment (DTE) communicate. 2) in telephony, it refers to the categorization of telephone subscribers according to specific type of telephone usage. Telephone service distinctions that include, for example, rate differences between individual and party lines, flat rate and message rate, and restricted and extended area service.

class of service display to attendant—shows the class of service for calls to the attendant console from inside extensions. The class of service shows which trunks or lines the extension is restricted from accessing.

class of service intercept—station is automatically routed to the attendant if it attempts to place a call that is not authorized by its class of service.

clear-forward/clear-back signal—signal transmitted from one end of a subscriber line or trunk, in the forward/backward direction, to indicate at the other end that the established connection should be disconnected. Also called *disconnect signal*.

clear to send—a control circuit that indicates to the data terminal equipment that data can or cannot be transmitted.

clipping—loss of initial or final parts of words or syllables due to operation of voice-actuated devices.

clock—in logic or transmission, repetitive, precisely timed signal used to control a synchronous process.

clocking—repetitive, regularly timed signals used to control synchronous transmissions.

closed circuit television (CCTV)—television transmission via direct link between two points, as opposed to broadcast transmission to many receiving locations.

closed user groups (CUG)—restricts access to and from one or more terminals to other members of the CUG (found on packet switched systems, E-mail, etc.).

cluster—two or more terminals connected to a single point or node.

cluster controller—a device that handles the remote communications processing for multiple (usually dumb) terminals or workstations.

CMOS—complementary metal-oxide semiconductor (logic circuit).

CO—see *central office*.

coaxial cable—a popular transmission medium consisting of one or more central wire conductors, surrounded by a dielectric insulator, and encased in either a wire mesh or extruded metal sheathing. There are many varieties, depending on the degree of EMI shielding afforded, voltages, and frequencies accommodated. Common CATV transmission cable typically supports RF frequencies from 50 to about 500M Hz.

Cobol—Common Business Oriented Language. Language used in programming computers.

COCOT—customer-owned, coin-operated telephone.

code—the conventions specifying how data may be represented in a particular system.

Glossary

code call access—feature that allows attendants and station users to dial an access code and a two- or three-digit called party code to activate signaling devices throughout a customer's premises with a coded signal corresponding to the called code (either audible or visible). The called or "paged" party can then be connected to the calling party by dialing a "meet-me" answering code from any station within the PBX system.

code character—set of conventional elements established by the code to enable the transmission of a written character (letter, figure, punctuation sign, arithmetical sign, etc.) or the control of a particular function (spacing, shift, line-feed, carriage return, phase corrections, etc.).

code restriction—denies selected station lines completion of dialed outgoing exchange network calls to selected exchange and area codes, both local and distant.

code ringing—alerting of telephone subscribers on multi-party lines by combinations of short and long rings that are different for each subscriber.

codec—coder-decoder device used to convert analog signals, such as speech, music, or television, to digital form for transmission over a digital medium, and back again to the original analog form. One is required at each end of the channel.

coin box—telephone, usually public, requiring insertion of coins before it can be used.

coin-value tones—tones produced in multislot coin telephones when different coins are deposited. The tones are detected and transmitted to the operator so that the correct amount can be checked. Also called *coin-denomination tones*.

collect call—see *reverse charge call*.

collision—overlapping transmissions that interfere with one another. Occurs when two or more devices attempt to transmit at or about the same instant.

Comité Consultatif International de Téléphonie et de Télégraphie (CCITT)—International Consultative Committee for Telephone and Telegraph; an advisory committee within the International Telecommunications Union (ITU) that recommends data communications standards.

command—a signal or group of signals which causes a computer to execute an operation or series of operations.

command-driven—programs requiring that the task to be performed be described in a special language with strict adherence to syntax. Compare to menu-driven.

common battery—DC power source in the exchange that supplies power to all subscriber stations and exchange office switching equipment.

common-battery signaling—method by which supervisory and telephone address information is sent to an exchange by depressing and releasing the switch on the cradle of the handset.

common bell—capability of an individual station ringer to respond to all incoming calls on all lines terminated on that instrument.

common business oriented language—see *Cobol*.

common carrier—an organization in the business of providing regulated telephone, telegraph, telex, and data communications services.

common channel interoffice signaling (CCIS)—an electronic means of signaling between any two switching systems independent of the voice path. The use of CCIS makes possible new customer services, versatile network features, more flexible call routing, and faster call connections.

common channel signaling system number 7—a CCITT-specified signaling protocol used in high-speed digital networks to provide communication between intelligent network nodes.

common control—automatic switching arrangement in which the control equipment necessary for the establishment of connections is shared, being associated with a given call only during the period required to accomplish the control function.

common control switching arrangement (CCSA)—switching facilities connected by the telephone company to incorporate tie line networks. Switching of the leased lines in the organization's network is accomplished by common control exchange switching equipment. All stations in the network can then dial one another regardless of distance and without using exchange facilities. They can also dial outside the network via local and/or foreign exchange lines.

communication—transmission of intelligence between points of origin and reception, without alteration of sequence or structure of the content.

communication control characters—in ASCII, a functional symbol intended to control or facilitate transmission over data networks.

communication line—any medium, such as a wire or a telephone circuit, that connects remote stations for the purpose of transmitting/receiving information.

communications channel—see *channel*.

Glossary

communications controller—dedicated computer with special processing capabilities for organizing and checking data and handling information traffic to and from many remote terminals or computers, including functions such as message switching.

communications processor—see *communications controller*.

communications satellite—earth satellite designed to act as a telecommunications radio relay and usually positioned in geosynchronous orbit 35,800 kilometers (23,000 miles) above the equator so that it appears from earth to be stationary in space.

community antenna television (CATV)—1) where television reception is poor, signals can be received at a selected site by sensitive, directional antennas and then transmitted to subscribers via a cable network. Additional channels, not normally available in that area, can also be transmitted. 2) data communications based on radio frequency (RF) transmission, generally using 75-ohm coaxial cable as the transmission medium. CATV offers multiple frequency-divided channels, allowing mixed transmissions to be carried simultaneously.

compandor—combination of a compressor at one point in a communications path for reducing the volume range of signals, followed by an expander at another point for restoring the original volume range. Designed to improve the ratio of the signal to the interference entering in the path between the compressor and expander.

compatibility—the characteristic of data processing equipment by which one machine may accept and process data prepared by another machine without conversion or code modification.

compression—the application of any of several techniques that reduce the number of bits required to represent information in data transmission or storage, therefore conserving bandwidth and/or memory.

compressor—electronic device that compresses the volume range of a signal.

computer database—a collection of facts in machine-readable form. Computer databases vary according to type (e.g., numeric, bibliographic, full-text) and to availability (proprietary or public).

computer network—an interconnection of two or more computer systems, terminals, and/or communications facilities.

computerized branch exchange (CBX)—see *private automatic branch exchange*.

COMSAT—Communications Satellite Corporation, a private U.S. company established by statute as the exclusive international satellite carrier and representing the U.S. in Intelsat.

concatenation—the linking of transmission channels (phone lines, coaxial cable, optical fiber) end-to-end.

concentrator—device that connects a number of circuits that are not all used at once to a smaller group of circuits for economical transmission. A telephone concentrator achieves the reduction with a circuit-switching mechanism.

conditioning—procedure to make transmission impairments of a circuit lie within certain specified limits and typically used on telephone lines leased for data transmission to improve the possible transmission speed. Two types are used: C conditioning and D conditioning. Also called *line conditioning*.

conductor—1) any equipment, such as a wire or cable, that can carry an electric current. 2) one wire of a pair of wires.

conduit—pipe or tubing through which cables are passed.

conf—conference.

conference call—call established among three or more stations in such a manner that each of the stations is able to carry on a communication with the others.

Conference of European Posts and Telecommunications (CEPT)—an organization formed by the European PTTs for the discussion of operational and tariff matters.

configuration—a group of machines interconnected and programmed to operate as a system.

connect time—the amount of time that a circuit, typically in a circuit-switched environment, is in use; see also, *holding time*.

cons—consultation.

console—the part of a computer used for communicating between the operator or maintenance engineer and the computer. A CRT terminal or a typewriter console are the most common.

console-less operation—allows internal stations to answer all incoming calls and process them to the proper stations.

consultation hold—incoming call is automatically placed on hold and the station user is given a PBX system dial tone. The station user proceeds to establish connection with another station. The original station can return to the incoming call.

Glossary

consultation hold, all calls—similar to the consultation hold facility, but not restricted to incoming calls.

contact—part of a switch designed to touch a similar contact to allow current to flow and to break this union to cause a current to cease.

contention—method of line control in which the terminals request to transmit. If the channel in question is free, transmission proceeds; if it is not free, the terminal will have to wait until it becomes free. A queue of contention requests can be built up by a computer, and this queue can either be organized in a prearranged sequence or in the sequence in which requests are made.

continuity check—in common-channel signaling, a test performed to check that a path exists for speech or data transmission.

contl—control.

control character—a character inserted into a datastream for signaling the receiving station to perform a function or to identify the structure of the message. Newer protocols are getting away from character-oriented control procedures into bit-oriented control procedures.

control mode—state that all terminals on a line must be in to allow line control actions or terminal selection to occur.

control procedure—the method of communicating information in an orderly way between stations on a data link. Synonymous with link discipline; see also *protocol*.

control signals—signals which pass between one part of communication system and another to control the system.

control, stored-program—means of system control using stored logic (software).

control, wired-program—means of system control using wired logic (hardware).

control station—the network station that supervises procedures such as polling, selecting, and recovery. It also is responsible for maintaining order in the event of any abnormal situation.

controlled station-to-station restriction—allows the attendant to inhibit station-to-station calling. When activated, attempted station calls are rerouted to attendant intercept.

control unit—the section of a computer that directs the sequence of operations, interrupts coded instructions, and sends the proper signals to other computer circuits to carry out the instructions. Also, an auxiliary component of a computer located behind or within the mainframe and

other component equipment, such as tape units, printers, and optical readers, for the purpose of controlling these components.

conversational mode—a mode of operation of a data processing system in which a sequence of alternating entries between a user and the system takes place in a manner comparable to a conversation between two persons.

conversion—the process of changing from one method to another; may refer to changing processing methods, data, or systems.

co-processor—an additional central logic unit which performs specific tasks while the main unit executes its primary tasks. Frequently, these chips are added to speed up mathematical tasks or perform I/O functions.

COS—see *class of service*.

country code—in direct distance dialing, a code characterizing a particular country. Codes corresponding to the world numbering plan start with a single digit that identifies a particular geographical area. This can be followed by one or two extra digits.

CPE—see *customer premises equipment*.

CPI—computer to PBX interface. Gateway providing 24 digital PCM channels at an aggregate rate of 1.544M bps. Developed by Northern Telecom. See also *DMI*.

CP/M—Control Program/Microcomputer; a family of operating systems developed by Digital Research, Inc.

cps—characters per second.

CPT—Conference of European Posts and Telecommunications. An organization formed by the European PTTs for the discussion of operational and tariff matters.

CPU—see *central processing unit*.

CR—see *carriage return*.

crash—a breakdown resulting from hardware or software malfunction.

CRC—see *cyclic redundancy check*.

critical path scheduling—a project planning and monitoring system used to check progress toward completion of the project by scheduling events, activities, milestones, etc.

cross-modulation—interference caused by two or more carriers in a transmission system interacting through nonlinearities in the system.

Glossary

cross section—signal transmission capacity of a transmission system, usually measured in terms of the number of two-way voice channels.

crossbar switch—switch having multiple vertical and horizontal paths and an electromagnetically operated mechanical means for interconnecting any one of the vertical paths within any one of the horizontal paths.

crosspoint—1) switching element in an exchange that can be mechanical or electronic. 2) two-state semiconductor switching device having a low transmission system impedance in one state and a very high one in the other.

crossstalk—interference or an unwanted signal from one transmission circuit detected on another, usually parallel, circuit.

crossstalk attenuation—extent to which a communications system resists crossstalk.

CRR—customer service record.

CRT—see *cathode-ray tube*.

CSMA—see *carrier sense multiple access*.

CSMA/CD—see *carrier sense multiple access with collision detection*.

CSR—Customer Service Record.

CSU—see *channel service unit*.

CTA—see *circuit terminating arrangements*.

CT2—second generation (digital) cordless telephone specification.

C/T—carrier-to-noise-temperature. Expressed in degrees Kelvin.

CUG—see *closed user groups*.

current loop—a transmission technique that recognizes current flows, rather than voltage levels. It has traditionally been used in teletypewriter networks incorporating batteries as the transmission power source.

cursor—a position indicator employed in CRT terminals to indicate a position in which data is to be entered.

cursor keys—keys that control the movement of the position indicator in the CRT; usually designated with arrows; see also *cursor*.

custom key set—specialized multibutton telephones designed expressly for a particular PBX. Unlike the locking buttons on normal key telephones, the buttons on a custom

key set are used to communicate with the system and are typically nonlocking buttons. Custom key set buttons can be arranged to activate specific features such as speed dialing and executive override as well as to select lines.

customer access line charges (CALC)—basic rates paid for lines from customer premises to central office.

customer administration center system (CACS)—allows the customer to administer station and electronic tandem switching features as well as obtain traffic measurements and recent circuit assurance data from one or more PBX switching locations. The user operates the CACS system via an interactive terminal. The following features can optionally be provided: station rearrangement and change, facilities administration and control, traffic data to customer, and facilities assurance reports.

customer administration panel (CAP)—provides a panel that is a simplified alternative to the customer administration center system. With this panel, the customer has plug-in access to a PBX for the purpose of performing administration of station features and/or electronic tandem switching capabilities.

customer information control system (CICS)—IBM communications monitor software system with database capabilities.

customer premises equipment (CPE)—terminal equipment, supplied by either the telephone common carrier or by a competitive supplier, that is connected to the telephone network.

customer service record (CSR)—detailed description of a subscriber's service and equipment charges.

cutover—physical changing of lines from one system to another, usually at the time of a new system installation.

CVSD—Continuous Variable Slope Delta Modulation. Speech encoding and digitizing technique that uses 1 bit to describe the change in the slope of the curve between 2 samples, rather than the absolute change between the samples. Sampling is usually done at 32K bps.

CW—call waiting.

CXR—see *carrier*.

cyclic redundancy check (CRC)—a powerful error detection technique. Using a polynomial, a series of two 8-bit block check characters are generated that represent the entire block of data. The block check characters are incorporated into the transmission frame, then checked at the receiving end.

Glossary

D

d conditioning—type of conditioning that controls harmonic distortion and signal-to-noise ratio so that they lie within specified limits.

DA—don't answer.

DAA—see *data access arrangement*.

DACS—see *digital access and cross-connect system*.

daisy chaining—the connection of multiple devices in a serial fashion. An advantage of daisy chaining is a savings in transmission facilities. A disadvantage is that if a device malfunctions, all of the devices daisy chained behind it are disabled.

DAMA—see *demand assigned multiple access*.

DASS—a public network standard for Digital Access Signaling System, developed by the United Kingdom telecommunications industry. The present version is DASS2, a message-based signaling system based on the ISO model, developed by British Telecom and its suppliers for multi-line integrated digital access to the public network.

data—units of information which can be precisely defined; raw facts and figures which are processed into information.

data access arrangement (DAA)—an interface device used to allow interconnect customer-owned data transmission equipment (DTE) to the telephone network; now generally integrated into such directly attached devices. DAAs were required in the 1970s; they are no longer mandatory.

data center—common designation for the computer-equipped, central location within an organization. The center processes and converts information to a desired form such as reports or other types of management information records.

data circuit—a communications facility that allows transmission of information in digital form.

data circuit terminating equipment (DCE)—in a communications link, equipment that is either part of the network, an access point to the network, a network node, or equipment at which a network circuit terminates. In the case of an RS-232-C connection, the modem is usually regarded as DCE while the user device is DTE, or data terminal equipment; in a CCITT X.25 connection, the network access and packet-switching node is viewed as the DCE.

data communications—1) the transmission, reception, and validation of data. 2) data transfer between a data source and a data destination via one or more data links, according to appropriate protocols by means of an electromagnetic transmission system.

data compression—a technique that saves storage space by eliminating gaps, empty fields, redundancies, or unnecessary information to shorten the length of records or blocks.

data concentration—collection of data at an intermediate point from several low- and medium-speed lines for retransmission across high-speed lines.

data conversion—the process of changing data from one form of representation to another.

data encryption standard (DES)—a cryptographic algorithm designed by the National Bureau of Standards (now the National Institute of Standards and Technology) to encipher and decipher data using a 64-bit key.

data line interface (DLI)—point at which a data line is connected to a telephone system.

data line privacy—critical system extension lines, used with such devices as facsimile machines and computer terminals, are very sensitive to extraneous noise. Data privacy prohibits activities that would insert tones on the station line while it is in use. Data lines can then be connected through the PBX without danger of losing data through interference.

data link—any serial data communications transmission path, generally between two adjacent nodes or devices and without any intermediate switching nodes.

data link control—1) procedures to ensure that both the sending and receiving devices agree on synchronization, error detection and recovery methods, and initialization and operation methods for point-to-point or multipoint configurations. 2) the second layer in the ISO Reference Model for Open Systems Interconnection.

data PBX—a switch that allows a user on an attached circuit to select from other circuits, usually one at a time and on a contention basis, for the purpose of establishing a through connection. A data PBX is distinguished from a PBX in that only digital transmission, and not analog, is supported.

data rate—the speed at which a channel carries data, measured in bits per second (bps).

data service unit (DSU)—simplified modem for the transmission of digital data over a private line, or for limited distance communications over the public switched telephone network (PSTN) where it is not necessary to comply with all the requirements for a high speed modem.

data set—infrequently used term for a modem.

data terminal equipment (DTE)—generally end-user devices, such as terminals and computers, that connect to DCE, which either generate or receive the data carried by

Glossary

the network. In RS-232-C connections, designation as either DTE or DCE determines signaling role in handshaking; in a CCITT X.25 interface, the device or equipment that manages the interface at the user premises. Compare with *DCE*.

data transmission—the process of transmitting information from one location to another by means of some form of communications media.

database—an organized compilation of computerized data maintained by a single organization or company. There are formalized (CODASYL) rules for the establishment of a database, as well as standardized terminology, and suggested means for controlling the data and access to it. The goal of a database is to eliminate redundant files and to assure that all data processed by the computer are current.

datagram—packet of data which can be delivered through a packet switched system without reference to previous packets addressed to the same destination.

Dataphone—an AT&T trademark. 1) an AT&T modem family. 2) an AT&T digital data service offering.

datel—data transmission services offered by European PTTs, using public switched telephone networks (PSTNs).

dB—see *decibel*.

DB/DC systems—database/data communications systems.

dBK—figure of merit (antenna system performance).

dBm—decibels relative to one milliwatt.

dBr—relative level in decibels.

DCE—see *data circuit terminating equipment*.

DCS—see *digital cross connect system*.

DDC—see *direct department calling*.

DDCMP—see *digital data communications message protocol*.

DDD—see *direct distance dialing*.

DDI—direct dialing in. See *direct inward dialing*.

DDS—see *digital data service*.

decibel (dB)—unit for measuring relative strength of a signal parameter such as power or voltage. The number of decibels is 10 times the logarithm (base 10) of the ratio of the power of two signals, or ratio of the power of one signal to a reference level.

dedicated—used exclusively for a single purpose or by a single subscriber.

dedicated attendant link—assures the availability of an intercom link for the attendant announcing incoming calls to the station within the system.

dedicated circuit—end-to-end communications line used exclusively by one organization.

delay—in communications, the wait time between two events, such as from when a signal is sent until it is received (see *propagation delay*, *response time*).

delay announcements—associated with Automatic Call Distribution capability, provides a recorded announcement to incoming calls that are delayed and placed in queue. The user can vary the time interval between the first and second delay announcements.

delay distortion—the change in a signal from the transmitting end to the receiving end resulting from the tendency of some frequency components within a channel to take longer to be propagated than others.

delimiter—in data communications, a character that separates and organizes elements of data.

delta modulation—method of representing a speech waveform (or other analog signal) in which successive bits represent increments of the waveform. The increment size is not necessarily constant.

deluxe queuing—allows station users, tie trunks and attendants, or attendant-assisted calls to be placed in a queue whenever all routes for completing a particular call are busy. The queue can be a Ringback Queue (RBQ), in which case a user goes on-hook and is called back when a trunk becomes available, or an Off-Hook Queue (OHQ), in which case the user remains off-hook and is connected to a trunk when it becomes available. Four types of specialized queuing arrangements are available; these are combinations of RBQ and OHQ with one or two queues.

demand assigned multiple access (DAMA)—allocation of communication satellite time to earth stations as the need arises.

demarc—demarcation point between carrier equipment and customer-premises equipment (CPE); usually a terminal block.

demodulation—the opposite of modulation; the conversion of a signal from analog to its original (e.g., digital) form.

derivation equipment—equipment used to produce narrow band facilities from a wider band facility. Commonly used on alternate voice/data (AVD) circuits to derive telegraph-grade lines from unused portions of a voice circuit.

Glossary

DES—see *data encryption standard*.

destination field—a field in a message header that contains the address of the station to which a message is being directed.

diagnostics—software routines used to check equipment malfunctions and to pinpoint faulty components.

dial—device that transmits a coded signal to actuate the exchange switching equipment according to the digit dialed.

dial call pickup—station user can dial a special code and thereby answer incoming calls ringing on any other station within his/her own predefined pickup group.

dial pulse (DP)—current interruption in the DC loop of a calling telephone produced by the breaking and making of the dial pulse contacts of a calling telephone when a digit is dialed.

dial selection of stations—indicates that the system accepts dialing of the number of the desired intercom station.

dial speed—the number of pulses that a rotary dial can transfer in a given amount of time, typically 10 pulses per second.

dial tone (DT)—signal sent to an operator or subscriber indicating that the switch is ready to receive dial pulses.

dial "0" trunks to attendant—special single-digit dialing service by any station to the attendant console or cord switchboard.

dial-up—the process of, or the equipment or facilities involved in, establishing a temporary connection via the switched telephone network.

dial-up line—a circuit that is established by a switched circuit connection; generally refers to the AT&T telephone network.

dictation access and control—allows station users to have dial access to centralized dictation equipment and to maintain telephone dial control of all normal dictation system features.

DID—see *direct inward dialing*.

digital—referring to communications procedures, techniques, and equipment whereby information is encoded as either binary "1" or "0"; the representation of information in discrete binary form, discontinuous in time, as opposed to the analog representation of information in variable, but continuous, waveforms.

digital access and cross-connect system (DACCS)—an AT&T term for a digital cross-connect system (DCS). See *digital cross-connect system*.

digital cross-connect system (DCS)—a computerized facility allowing DS1 lines (1.544M bps) to be remapped electronically at the DS0 level (64K bps) meaning that DS0 channels can be individually re-routed and reconfigured into different DS1 lines.

digital data communications message protocol (DDCMP)—a synchronous protocol developed by Digital Equipment Corporation.

digital data service (DDS)—a digital transmission service supporting speeds up to 56K bps.

digital loopback—technique for testing the digital processing circuitry of a communications device. It can be initiated locally, or remotely via a telecommunications circuit; the device being tested will echo back a received test message, after first decoding and then reencoding it, the results of which are compared with the original message—compare with *analog loopback*.

digital network—network incorporating both digital switching and digital transmission.

Digital Network Architecture (DNA)—the network architecture of Digital Equipment Corporation.

digital signal—discrete or discontinuous signal; one whose various states are discrete intervals apart.

digital speech interpolation (DSI)—when speech is digitized, it can be cut into slices such that no bits are transmitted when a person is silent. As soon as speech begins, bits flow again.

digital switching—the process of establishing and maintaining a connection, under stored program control, where binary-encoded information is routed between an input and an output port. Generally, a virtual through circuit is derived from a series of time slots (time-division multiplexing), which is more efficient than requiring dedicated circuits for the period of time that connections are set up.

DIP—dual in-line package.

DIP switch—a dual in-line switch; allows the user to set current paths on or off; they are frequently used to reconfigure microcomputer components and peripherals.

direct attendant signaling lines—station user (momentarily) activates the pushbutton to signal the attendant, and when time permits, the attendant rings the station and verbally seeks instructions.

Glossary

direct call pickup—station user is able to answer calls ringing on any other station within the PBX system by dialing the unique answer code of that station to be answered. If the call has already been answered, the station user who dials the answer code can be added (in conference) to the connection.

direct department calling (DDC)—incoming calls on a specific trunk or group of trunks are routed to specific stations or groups of stations.

direct distance dialing (DDD)—telephone exchange service that enables the telephone user to call long distances without operator assistance.

direct-in lines—allow direct termination of separate exchange lines to station instruments, bypassing the attendant console.

direct inward dialing (DID)—incoming calls from the exchange network can be completed to specific station lines without attendant assistance. Also called *direct dialing in (DDI)*.

direct inward system access (DISA)—feature that allows an outside caller the ability to dial directly into a PBX system, without attendant intervention, and gain complete access to PBX system facilities and outgoing trunk circuits and tie line circuits.

direct outward dialing (DOD)—allows a PBX, Centrex, or hybrid system user to gain access to the exchange network without the assistance of the attendant.

direct station selection (DSS)—ability to call any station within a PBX by means of a single pushbutton associated with that station number. A DSS panel is usually part of an attendant console.

direct trunk group selection—allows the attendant to access an outgoing trunk from a particular group by pushing a single button associated with the group rather than dialing an access code.

directory number—the full complement of digits (or letters and figures) associated with the name of a subscriber in the directory.

DISA—see *direct inward system access*.

disconnect signal—see *clear-forward clear-back signal*.

discriminating ringing—different types of station ringing are provided to give audible indication of internal or external incoming calls and other special calls.

display systems protocol (DSP)—IBM protocol which allows 3270 BSL control units, printers, terminals, etc., to interface with PDN.

distinctive dial tones—internal calls, external calls, and internal calls originated with a caller on hold are provided with different dial tones.

distinctive ringing tones—different ringing tones provided for an extension user to indicate whether the call is internal or external.

distortion—the modification of the waveform or shape of a signal caused by outside interference or by imperfections of the transmission system. Most forms of distortion are the result of the varying responses of the transmission system to the different frequency components of the transmission signal.

distributed data processing (DDP)—the use of computer systems or intelligent terminals within an organization to perform data processing and/or storage functions. The devices may or may not be interconnected via transmission facilities.

Distributed Systems Architecture (DSA)—the network architecture developed by Honeywell.

Distributed Systems Network (DSN)—the network architecture developed by Hewlett-Packard.

distribution frame—structure (typically wall-mounted) for terminating telephone wiring, usually the permanent wires from, or at, the telephone exchange, where cross-connections are readily made to extensions. Also called *distribution block*.

diversity—provision of more than one communications channel to improve the reliability of a service.

divestiture—the breakup of AT&T in the U.S., mandated in 1984 by the federal courts based on an antitrust accord reached between AT&T and the U.S. Department of Justice. The decision included the separation of 22 AT&T-owned local Bell Operating Companies (BOCs) into seven independent Regional Bell Holding Companies (RBHCs).

DMI—Digital Multiplexed Interface. Gateway providing 23 digital PCM channels + 1 signaling channel at an aggregate rate of 1.544M bps. Developed by AT&T. See also *CPI*.

DNA—see *Digital Network Architecture*.

do not disturb—user can instruct the system to cancel all ringing at the station.

DOD—see *direct outward dialing*.

DOMSAT—domestic communications satellite.

don't answer recall—for internal or outside calls, a station user who experiences a "don't answer" condition at the

Glossary

called station can have the call automatically retried at a later time by the use of a preassigned single- or double-digit code.

double camp-on indication—station attempting to camp-on to another station that is already being camped-on receives a distinctive audible signal and can be denied the ability to camp-on.

downlink—the portion of a satellite circuit extending from the satellite to the earth. Compare to *uplink*.

download—the process of loading software into the nodes of a network from one node over the network media.

downtime—the total time a system is out of service due to equipment failure.

DP—1) dial pulse (signaling). 2) distribution point.

DPC—data processing center.

DPNSS—a digital private network signaling system standard developed by British Telecom and other United Kingdom PBX suppliers.

drop—a connection point between a communicating device and a communications network.

drop, subscriber's—line from a telephone cable to a subscriber's building.

DS-C1—digital signal level 1C; telephony term describing a 3.152M bps digital signal carried on a T1 C facility.

DS0—digital signal level 0; telephony term for a 64K bps standard digital telecommunications signal or channel.

DS1—digital signal level 1, a digital transmission format in which 24 voice channels are multiplexed into one 1.544M bps (U.S.) or 2.108M bps (Europe) T1 digital channel.

DS2—digital signal level 2; telephony term describing the 6.312M bps digital signal carried on a T2 facility.

DS3—digital signal level 3; telephony term describing the 44M bps digital signal carried on a T3 facility.

DS4—digital signal level 4; telephony term describing the 273M bps digital signal carried on a T4 facility.

DSA—see *Distributed Systems Architecture*.

DSI—see *digital speech interpolation*.

DSN—see *distributed systems network*.

DSP—see *display systems protocol*.

DSS—see *direct station selection*.

DSU—see *data service unit*.

DT—see *dial tone*.

DTE—see *data terminal equipment*.

DTMF—see *dual tone multifrequency signaling*.

dual in-line package (DIP)—method of packaging electronic components for mounting on printed circuit boards.

dual tone multifrequency signaling (DTMF)—basis for operation of pushbutton telephone sets. A method of signaling in which a matrix combination of two frequencies, each from a group of four, is used to transmit numerical address information. The two groups of four frequencies are 697 Hz, 770 Hz, 852 Hz, and 941 Hz, and 1209 Hz, 1336 Hz, 1477 Hz, and 1633 Hz.

dumb terminal—a conversational display terminal with limited editing capabilities.

dump—to transfer all information from a record to another storage medium; e.g., copying from memory to a printer.

duplex—synonymous with full-duplex.

duplex circuit—circuit used for transmission in both directions at the same time. It can be called "full-duplex" to distinguish it from "half-duplex." See *full-duplex* and *half-duplex*.

duplex signaling (DX)—signaling system that occupies the same cable pair as the voice path yet does not require filters.

duplex transmission—simultaneous, two-way, independent transmission. Also called *full-duplex transmission*.

DX—see *duplex signaling*.

E

E—see *erlang*.

E&M signaling—signaling arrangement that uses separate paths for signaling and voice signals. The M lead (derived from "mouth") transmits ground or battery to the distant end of the circuit, while incoming signals are received as either a grounded or open condition on the E (derived from "ear") lead. E&M is the traditional, inband, send and receive signaling method used by the North American Public Switched Telephone network. Some PBXs use E&M signaling, also. Most telcos are replacing E&M with computer-controlled CCIS (out-of-band) signaling, which is required for ISDN implementation.

EAROM—electrically alterable read-only memory.

Glossary

earth station—ground-based equipment used to communicate via satellites. See also *ground station*.

EAS—see *extended area service*.

EAX—electronic automatic exchange.

EBCDIC (extended binary coded decimal interchange code)—a coded character set comprising 8-bit coded characters, developed by IBM and a de facto industry standard.

ECC—see *error correcting code*.

echo—wave that has been reflected or otherwise returned with sufficient magnitude and delay for it to be perceived as a wave distinct from that directly transmitted.

echo cancellation—technique being used in new higher speed modems that allows for the isolation and filtering of unwanted signal energy resulting from echoes from the main transmitted signal.

echo canceller—a device used to reduce or eliminate echo. It operates by placing a signal that is equal and opposite to the echo signal on the return transmission path.

echo check—method of checking data transmission accuracy whereby the received data are returned to the sending end for comparison with the original data.

echo return loss (ERL)—attenuation of echo currents in one direction caused by telephone circuits operating in the other direction.

echo suppressor—a mechanism used to suppress interference on long-distance analog connections; it suppresses the transmission path opposite in direction to the one being used. This feature is necessary for voice transmission but often interferes with data transmission.

echoplex—pertaining or referring to the printing of keyboarded characters on return of the signal from the other end of the line, using full-duplex transmission, to assure that the data was received correctly at the other end.

ECL—emitter-coupled logic. Type of ALU technology.

ECMA—European Computer Manufacturers Association.

ECS—European fixed-service satellite system.

EDC—see *error detecting code* and *error correcting code*.

EDI—see *electronic data interchange*.

edit—the modification of data, format, and code conversion and the application of processes such as zero suppression.

EEPROM—electrically erasable programmable read-only memory. A memory that can be electronically programmed and erased, but which does not require a power source to retain data.

EIA—Electronic Industries Association.

EIA interface—a standardized set of signal characteristics (time duration, voltage, and current) specified by the Electronic Industries Association.

EIRP—effective isotropic radiated power. The combination of transmitted power and antenna gain.

EKTS—electronic key telephone system.

electromagnetic interference (EMI)—the energy given off by electronic circuits and picked up by other circuits; based on type of device and operating frequency, EMI can be reduced by shielding; minimum acceptable levels are detailed by the FCC.

electromagnetic spectrum—entire range of wavelengths or frequencies of electromagnetic radiation extending from gamma rays to the longest radio wave and including visible light. See also *infrared* and *radio frequency*.

electromechanical ringing—bell or buzzer provided in the station instrument to give audible incoming call indication.

electronic automatic exchange (EAX)—term used by the General Telephone Company for electronic exchange equipment.

electronic data interchange (EDI)—paperless electronic exchange of trading documents, such as invoices and orders.

electronic mail (E-mail)—the transmission of letters, messages, and memos from one computer to another without using printed copies.

electronic switching system (ESS)—system using computer-like operations to switch telephone calls.

electronic tandem networking—operation of two or more switching systems in parallel.

elevation—the angle between the horizon and an antenna's beam.

emergency access—emergency alarm system integrated into the PBX switching system that can be implemented by uniquely ringing all station instruments to indicate an emergency condition.

emergency dialing—allows all system users the ability to use preassigned two- or three-digit numbers that translate to the outside exchange numbers for local police, local fire, ambulance service, etc.

Glossary

EMI—see *electromagnetic interference*.

emulate—to imitate one system with another, so that the imitating system accepts the same data, executes the same computer programs, and achieves the same results as the imitated system.

EN—equipment number.

encryption—the process of systematically encoding a bit stream before transmission so that an unauthorized party cannot decipher it.

end office—class 5 central office, at which subscriber loops terminate.

end-to-end test—as used by the Bell System, a method of exercising a communications link; it requires Bell maintenance personnel at each end of the circuit.

energy communications—provides the ability for the PBX to communicate with energy-consuming devices throughout a hotel or business complex. Audio signals are sent over the telephone wiring to power unit consumption devices in the building. This feature is designed to use existing telephone wiring wherever possible. A separate interface unit is required for each power unit. Specific operating modes include:

- Energy consumption and demand monitoring,
- Guestroom cycling function,
- Individual load cycling function,
- Peak demand shedding function, and
- Vacant room energy function.

engaged tone—see *busy signal*.

enhanced private switched communications service (EPSCS)—AT&T Communications service providing a private interLATA electronic tandem tie-line network linking several of a customer's locations together.

ENQ (Enquiry)—in data communications, a request for a response from another terminal; it is used to obtain identification and/or an indication of the other station's status.

envelope delay—an analog line impairment that conditioning is meant to overcome. A variation of signal delay with frequency across the data channel bandwidth.

EOM—end-of-message (indicator).

EOT—1) end of transmission. 2) end of tape.

EPABX—electronic private automatic branch exchange. See *exchange, private automatic branch*.

EPROM—erasable and programmable read-only memory.

EPSCS—see *enhanced private switched communication service*.

equal access—Department of Justice ruling (effective 9/84) that requires all RBHCs with ESS systems using SPC technology, and serving a market of at least 10,000 access lines, to offer the same quality of connection at the same rates to all common carriers. Under this arrangement, subscribers choose a primary long-distance carrier that they access by dialing "1" + area code + telephone number. Other long-distance carriers are reached by dialing "10" plus a three-digit access code unique to each carrier, then dialing "1" + area code + telephone number.

equalization—the introduction of components to an analog circuit by a modem to compensate for signal attenuation and delay distortion. Generally, the higher the transmission rate, the greater the need for equalization.

ergonomics—a discipline that promotes the consideration of human factors in the design of the working environment and its components (heat, light, sound, equipment, etc.).

ERL—see *echo return loss*.

Erlang (E)—unit of traffic intensity. One erlang is the intensity at which one traffic path would be continuously occupied, e.g., 1 call-hour per hour, 1 call-minute per minute, etc., generally referred to as 36 hundred call seconds (ccs).

Erlang B—traffic engineering formula used when traffic is random and there is no queuing. Erlang B assumes that blocked callers either automatically use another route, or blocked calls disappear entirely.

Erlang C—traffic engineering formula used when traffic is random and queuing is provided. Erlang C assumes that all callers will wait indefinitely until a line becomes available.

ERMES—program sponsored by the Commission of the European Communities for a pan-European radio(-paging) message system.

error—in data communications, any unwanted change in the original contents of a transmission.

error burst—a concentration of errors within a short period of time as compared with the average incidence of errors. Retransmission is the normal correction procedure in the event of an error burst.

error control—a process of handling errors, which includes the detection and correction of errors.

error correcting code—code incorporating sufficient additional signaling elements to enable the nature of some or all of the errors to be indicated and corrected entirely at the receiving end.

Glossary

error detecting code—code in which each telegraph or data signal conforms to specific rules of construction so that departures from this construction in the received signals can be automatically detected.

error rate—the ratio of the amount of data incorrectly received to the total amount of data transmitted.

ESA—European Space Agency.

ESPRIT—European Strategic Program for Research and Development in Information Technologies.

ESS—see *electronic switching system*.

ETACS—extended TACS. See *TACS*.

Ethernet—a local area data network, developed by Xerox Corporation and supported by Intel Corp., Digital Equipment Corp., and Hewlett-Packard.

ETN—see *electronic tandem networking*.

ETSI—European Telecommunications Standards Institute.

ETX—end-of-text (indicator).

even parity check (odd parity check)—tests whether the number of digits in a group of binary digits is even (even parity check) or odd (odd parity check).

EUTELSAT—European Telecommunications Satellite Organization. Operates the Eutelsat services of satellites.

exchange—assembly of equipment in a communications system that controls the connection of incoming and outgoing lines and includes the necessary signaling and supervisory functions. Different exchanges, or switches, can be co-sited to perform different functions, e.g., local exchange, trunk exchange, etc. See *class of exchange*. Also known as *central office*.

exchange area—a calling area served by a local telco central office switch, may comprise one or more NXX codes.

exchange line socket—2- or 4-wire socket providing physical access to the PSTN.

exchange, private (PX)—exchange serving a particular organization and having no means of connection with a public exchange.

exchange, private automatic (PAX)—dial telephone exchange that provides private telephone service to an organization and that does not allow calls to be transmitted to or from the public telephone network.

exchange, private automatic branch (PABX)—private automatic telephone exchange that provides for the transmission of calls internally and to and from the public telephone network.

exchange, private branch (PBX)—see *private automatic branch exchange (PABX)*.

exchange, private digital—see *PDX*.

exciter—the transmitter components that modulate and drive a signal. Used in uplinks.

exclusive hold—only the station that has placed a line circuit on hold is capable of breaking the hold condition and reestablishing conversation.

executive busy override—allows preselected stations to dial a single digit and gain access to the conversation taking place upon encountering a busy signal.

expander—transducer that, for a given amplitude range of input voltages, produces a larger range of output voltages.

extended area service (EAS)—option whereby the telephone subscriber can pay a higher flat rate in order to obtain wider geographical coverage without additional per-call charges.

extended binary-coded decimal interchange code (EBCDIC)—a coded character set consisting of 8-bit coded characters. EBCDIC is the usual code generated by synchronous IBM devices.

extension telephone—additional telephone set on the same line but at a different location than the main station.

extn—extension.

F

facilities administration and control—allows customer to administer the assignment of parameters that determine user calling privileges, such as restriction levels and authorization codes. Manual control (override) of time of day routing is provided. Activation and deactivation of trunk group queues are also provided.

facilities assurance reports—allows customer to obtain an audit trail of the data generated by the automatic circuit assurance feature. The audit trail indicates the identity of the trunk circuit, time of referral, date, nature of the referral (e.g., short-holding-time failure or long-holding-time failure), and whether a test was performed in response to the referral.

facilities restriction level—parameter associated with each authorization code, each station at a PBX, each incoming tie trunk group from subtending PBXs, and attendants as a whole, which determines both the types of calls and types of

Glossary

facilities within the privileges of the associated user. These restriction levels are used in routing calls via Automatic Route Selection (ARS) and Uniform Numbering/Automatic Alternate Routing (UN/AAR) features. Each route in each route pattern has a minimum restriction level required to use it. This feature replaces the code restriction level that can be used with ARS. When the restriction level of the calling party is transmitted over an intertandem tie trunk to a distant PBX, it is called a traveling class mark (TCM).

facility—1) any or all of the physical elements of a plan used to provide communications services. 2) a component of an operating system. 3) transmission path between two or more points, provided by a common carrier.

facsimile—system for the transmission of images, usually over the public telephone network. The image is scanned at the transmitter, reconstructed at the receiving station, and duplicated on some form of paper.

fading—a phenomenon, generally of microwave or radio transmission, where atmospheric, electromagnetic, or gravitational influences cause a signal to be deflected or diverted away from the target receiver. The reduction in intensity of the power of a received signal.

fail softly—when a piece of equipment fails, the programs let the system fall back to a degraded mode of operation rather than let it fail completely.

far-end crosstalk—crosstalk that travels along a circuit in the same direction as the signals in that circuit. Compare with *near-end crosstalk*.

fault—a condition that causes any physical component of a system to fail to perform in acceptable fashion.

fault tolerance—the ability of a program or system to operate properly even if a failure occurs.

fax—see *facsimile*.

FCC—see *Federal Communication Commission*.

FDDI—see *fiber distributed data interface*.

FDM—see *frequency-division multiplexing*.

FDMA—see *frequency-division multiple access*.

FD or FDX—see *full-duplex*.

feature group—under the provisions of Equal Access, three types of local-line access must presently be provided by the BOCs to long-distance carriers:

—Feature Group A access, which amounts to dialing the carrier's 7-digit access number, as well as an account code

and a 10-digit long-distance number (you could use an automatic dialer for this purpose).

—Feature Group B access, which is a trunk-side connection accessed by dialing 950-XXXX plus a 10-digit long-distance number. Feature Group B provides better quality but is more expensive.

—Feature Group D access, which is similar to FGB, permits 1 + 10-digit dialing of the primary long-distance carrier and can be considered the most complete form of access.

Federal Communications Commission (FCC)—Washington, DC regulatory agency established by the Communications Act of 1934, charged with regulating all electrical and radio communications in the U.S.

federal telecommunications system—a government communications system administered by GSA; it covers 50 states, plus Puerto Rico and the Virgin Islands, and provides services for voice, teletypewriter, facsimile, and data transmission.

feedback—the return of part of the output of a machine, process, or system, to the computer as input from another phase; also refers to system messages that keep a user informed of system activities during a process.

feed horn—a small antenna at the focal point of a reflector antenna that gathers the signal.

FEP—see *front-end processor*.

fiber distributed data interface (FDDI)—an ANSI standard specifying a packet switched LAN-to-LAN backbone for transporting data at high throughput rates over a variety of multimode fibers.

fiber optic waveguides—thin filaments of glass through which a light beam can be transmitted for long distances by means of multiple internal reflections. Occasionally other transparent materials, such as plastic, are used.

fiber optics—a technology that uses light as a digital information carrier. Fiber optic cables (light guides) are a direct replacement for conventional coaxial cables and wire pairs. The glass-based transmission facilities occupy far less physical volume for an equivalent transmission capacity, which is a major advantage in crowded underground ducts. The fibers are immune to electrical interference. Also called *lightwave communications*.

FIFO—see *first-in- first-out*.

figures shift—physical shift in a teletypewriter (specifically Telex) that enables the printing of numbers, symbols, uppercase characters, etc.

Glossary

file server—in local networks, a station dedicated to providing file and mass data storage services to the other stations on the network.

file transfer protocol (FTP)—a service providing a family of commands for performing file and directory operations over the network, such as append, rename and delete files, list change, make and remove directories, check status, toggle switches, and ask for help.

filter—circuit designed to transmit signals of frequencies within one or more frequency bands and to attenuate signals of other frequencies.

firmware—permanent or semipermanent control coding implemented at a micro-instruction level for an application program, instruction set, operating routine, or similar user-oriented function.

first-in, first-out (FIFO)—refers to a method of coordinating the sequential flow of data through a buffer.

fixed night service—routes incoming exchange calls to pre-selected stations within a PBX system when the attendant is not on duty.

flat rate—fixed payment for service within a defined area, independent of use, with an additional charge for each call outside the area.

flexible intercept—allows completely flexible and instantly changeable intercept conditions to the attendant, such as unassigned number, temporary disconnect, etc.

flexible line ringing—indicates the ability to arrange each station within the system with complete flexibility in regard to the ringing on incoming outside calls.

flexible night selection—allows the attendant to “set up” night connections in accordance with day-to-day requirements.

flexible numbering of stations—station numbers can be assigned to lines at installation based on customer specifications and can easily be changed thereafter. Some PBXs allow the user to perform the modifications while others require the attention of a service person. This feature is frequently used in hotels where the extension number is the same as the room number.

flexible release—ability of the switching system to effect connection release either by the calling party or by the first party to hang up.

flexible route selection (FRS)—used with Centrex service to achieve least-cost routing.

flow control—the use of buffering and other mechanisms, such as controls that turn a device on and off, to prevent data loss during transmission.

FM—see *frequency modulation*.

footprint—1) the space a device occupies on a desk or work surface. 2) in satellite communications, the footprint is the precise area of the earth in which the signal can be received. A given satellite can have a footprint of up to several thousand miles in width.

foreign exchange (FX)—connects a customer's location to a remote exchange. This service provides the equivalent of local service from the distant exchange.

Fortran—FORMula TRANslation. Language used in programming computers.

forward channel—the communications path carrying voice or data from the call initiator to the called party.

forward error correction—a technique for correcting errors incurred in transmission over a communications channel by the receiver, without requiring the retransmission of any information by the transmitter. Typically involves a convolution of the transmitted bits and the appending of extra bits, using a common algorithm by both the receiver and transmitter.

FOTS—fiber optic transmission systems. See *fiber optic waveguides*.

four-wire channel—a circuit containing two pairs of wire (or their logical equivalent) for simultaneous (i.e., full-duplex) two-way transmission. Contrast with *two-wire channel*.

four-wire circuit—a circuit in which two pairs of conductors, one for the inbound channel and one for the outbound channel, are connected to the station equipment. Contrast with *two-wire circuit*.

frame—1) in high-level data-link control (HDLC), the sequence of contiguous bits bracketed by and including opening and closing flag (01111110) sequences. 2) in data transmission, the sequence of contiguous bits bracketed by and including beginning and ending flag sequences. 3) a set of consecutive digit time slots in which the position of each time slot can be identified by reference to a frame alignment signal. 4) in a time-division multiplexing (TDM) system, a repetitive group of signals resulting from a signal sampling of all channels, including any additional signals for synchronizing and other required system information.

frame-grabber—device that can seize and record a single frame of video information out of a sequence of many frames.

Glossary

framing—a control procedure used with multiplexed digital channels, such as T1 carriers, where bits are inserted so that the receiver can identify the time slots that are allocated to each subchannel; framing bits may also carry alarm signals indicating specific alarm conditions.

frequency—an expression of how frequently a periodic (repetitious) wave form or signal regenerates itself at a given amplitude. It can be expressed in hertz (Hz), kilohertz (KHz), megahertz (MHz), etc.

frequency bands—frequency bands are defined arbitrarily as follows:

Range (MHz)—	Name—
0.03-0.3	—Low frequency (LF)
0.3-3.0	—Medium frequency (MF)
3-30	—High frequency (HF)
30-300	—Very high frequency (VHF)
300-3000	—Ultra high frequency (UHF)
3000-30,000	—Super high frequency (SHF) (microwave)
30,000-300,000	—Extremely high frequency (EHF) (millimeterwave)

frequency coordination—international procedure to prevent interference between new and existing radio communications services. Under International Telecommunication Union (ITU) Radio Regulations, potential operators must consult countries and administrations whose services might be affected.

frequency-division multiple access (FDMA)—communicating devices at different locations sharing a multipoint or broadcast channel by means of a technique that allocates different frequencies to different users.

frequency-division multiplexing (FDM)—a technique of dividing a single communications line into several data paths of different frequencies, each supporting an independent information stream. Contrast with *time-division multiplexing*.

frequency modulation (FM)—one of three ways of modifying a sine wave signal to make it carry information. The sine wave or “carrier” has its frequency modified in accordance with the information to be transmitted. The frequency function of the modulated wave may be continuous or discontinuous.

frequency shift keying (FSK)—a method of modulation that uses two different frequencies to distinguish between a mark (digital 1) and a space (digital 0) when transmitting on an analog line. Used in modems operating at 1200 bps or slower.

fresnel zone—zone of clearance used in calculations during the setting up of a microwave link. Although not directly between the transmitter and receiver, an object near the line-of-sight path can interfere with the propagated signal.

front-end processor (FEP)—a dedicated communications system that intercepts and handles activity for the host. It may perform line control, message handling, code conversion, error control, and such applications functions as control and operation of special-purpose terminals; see also *communications controllers*.

FRS—see *flexible route selection*.

FSK—see *frequency shift keying*.

FTAM—file transfer access management protocol.

full-duplex (FD, FDX)—refers to a communications system or equipment capable of transmission simultaneously in both directions. Contrast with *half-duplex*.

full text database—this type of source database contains complete textual records of primary sources. These sources include newspapers, specifications, court decisions, journals, etc.

fully restricted stations—feature that denies selected station lines the ability to place or receive any but internal station-to-station calls.

functional split—division within an automatic call distributor (ACD) that allows incoming calls to be directed from a specific group of trunks to a specific group of agents.

functional test—test carried out under normal working conditions to verify that a circuit or a particular part of the equipment functions correctly.

FX—see *foreign exchange*.

G

gain—denotes an increase in signal power in transmission from one point to another, usually expressed in dB.

GAP—an ad hoc working group, Groupe d'Analysis et Prevision.

GaAs FETs—Gallium Arsenide Field Effect Transistors—a form of generating low-power but reliable microwave energy up to 50 Hz.

gate assignments—used in context of ACD equipment. Gates are made up of trunks that require similar agent processing. Individual agents can be reassigned from one gate to another gate by the customer via the supervisory control and display station. Also called *splits*.

Glossary

gateway—a conceptual or logical network station that serves to interconnect two otherwise incompatible networks, network nodes, subnetworks, or devices. Gateways perform a protocol-conversion operation across a wide spectrum of communications functions or layers.

geostationary satellite—satellite in orbit about 35,800 kilometers (23,000 miles) above the equator, remaining vertically above the same point and hence appearing to be stationary. Such a satellite can be used as part of a permanent communications system.

geosynchronous orbit—the position where communications satellites will remain in orbit over the same location above the earth's equator. It is the area, about 23,000 miles above the earth's surface, where a satellite's orbit velocity matches the rotation of the earth, causing it to remain stationary relative to a point on the earth.

GHz—gigahertz (one billion cycles per second).

gigabyte—one billion bytes.

gigahertz (GHz)—a frequency unit equal to one billion cycles per second.

global satellite communications system—international commercial communications system established by Intelsat, consisting of geostationary satellites above the Atlantic, Pacific, and Indian oceans owned by Intelsat and earth stations belonging to the individual countries.

GOSIP—U.S. and U.K. government OSI procurement specification.

grade of service—quality of telephone service provided by a system described in terms of the probability that a call will encounter a busy signal during the busiest hour of the day.

grd—ground.

ground circuit—1) circuit in which energy is carried one way over a metallic path and returned through the earth. 2) circuit connected to earth at one or more points.

ground start—telephony term describing a signaling method whereby one station detects that a circuit is grounded at the other end.

ground station—assemblage of communications equipment, including signal generator, transmitter, receiver, and antenna, that receives (and usually transmits) signals to/from a communications satellite. Also called *earth station*.

group band modem—modem designed for CCITT V.35, V.36, or V.37 standard for transmission over group band (60-108 KHz) circuits.

group call—special type of station hunting requiring a special access number that will allow a call to the special access number and ring the first available number in that group.

GAPS—made up of industrial, Member State, and Commission representatives in the European Community for the study and agreement of common approaches to various communication issues. Reports by the group have covered ISDN, broadband technologies and networks, mobile communications, and Open Network Architecture (ONA).

grps—groups.

Groupe Speciale Mobile (GSM)—formed by the CEPT to define a pan-European digital cellular mobile standard.

gunn diode—named for IBM scientist John Gunn (1963). Until recently, the standard form for generating microwave energy. Produces only a single fixed frequency.

H

half-duplex (HD or HDX)—a circuit designed for transmission alternately in either direction but not both directions simultaneously. Contrast with *full-duplex*.

hamming code—code using redundant bits to detect errors in data (transmission errors).

handoff—transfer of duplex signaling as a mobile terminal passes to an adjacent cell in a cellular radio network.

hands-free station—capability of the station user to have speakerphone operation on all calls.

handset—portion of the telephone containing the transmitter and receiver, which is handled when the telephone is used.

handshaking—exchange of predetermined signals for control when a connection is established between two modems or other devices.

hard wired—1) referring to a communications link, whether remote phone line or local cable, that permanently connects two nodes, stations, or devices. 2) descriptive of electronic circuitry that performs fixed logical operations by virtue of unalterable circuit layout, rather than under computer or stored-program control.

harmonic distortion—a waveform distortion, usually caused by the nonlinear frequency response of a transmission.

HD or HDX—see *half-duplex*.

HDLC—see *high-level data-link control*. Bit-oriented communication protocol developed by the ISO.

Glossary

header—the initial portion of a message, which contains any information and control codes that are not part of the text (e.g., routine, priority, message type, destination addressee, and time of origination).

header record—data containing common, constant, or identifying information for a group of records that follow.

headset—operator or attendant telephone set that consists of a telephone transmitter, a receiver, and cord and plug, designed to leave the operator's hands free.

headset/recorder jack—allows a headset or input plug from a recorder to be connected to the talk circuit on the station instrument.

hertz (Hz)—unit of electromagnetic frequency equal to one cycle per second.

heterogeneous (computer) network—a system of different host computers, such as those of different manufacturers. Compare with *homogeneous network*.

heuristic—exploratory methods of problem solving in which solutions are arrived at by an interactive, self-learning method.

hexadecimal—referring to a number system with 16 members represented by 0 through 9 followed by A through F. Used to identify the 16 possible bit patterns of a half-byte; two hex digits represent one byte. Synonymous with hex.

HF—see *high frequency*.

hierarchical (computer) network—a system in which processing and control functions are performed at several levels by computers specially suited in capability for the functions performed.

high frequency (HF)—portion of the electromagnetic spectrum, typically used in short-wave radio applications; frequencies approximately in the 3 MHz to 30 MHz range.

high-level data-link control (HDLC)—a bit-oriented protocol developed by the ISO.

high-usage (HU) route—direct trunks provided, where traffic volume warrants, to bypass a part of the DDD switching network. Also called *high-usage trunk*.

histogram—a graph of contiguous vertical bars representing a frequency distribution in which the groups or classes of items are marked at equal intervals in ascending order on the x axis, and the number of items in each class is indicated by a horizontal line segment drawn above the x axis at a height equal to the number of items in the class.

hold—feature whereby a station user can remain connected to a line while not off-hook to the line. PBXs have provided

a new range of hold features that are more easily implemented by means of their programmed intelligence than through mechanical arrangements.

holding time—length of time a communications channel is in use for each transmission. Includes both message time and operating time. Also called *connect time*.

homogeneous (computer) network—a system of similar computers such as those of one model by the same manufacturer; contrast with *Heterogeneous Network*.

hookswitch—see *switchhook*.

hoot and holler circuit—four-wire, private line circuit using manual or automatic voice signaling or ringdown to transmit voice information among dispersed sites, such as branch offices. Used primarily in financial institutions to provide timely communication of investment information among traders. Derived from the lively conversational mode of the investment trading environment.

host computer—the primary or controlling system in a multiple computer network operation.

host interface—the link between a communications processor or network and a host computer.

hot line—line serving two telephone sets exclusively, on which one set will ring immediately when the receiver of the other set is lifted.

hot standby—alternate equipment ready to take over an operation quickly if the equipment on which it is being performed fails.

hotel/motel console—usually located at the hotel front desk, provides room status information as well as outgoing call records for guest stations. A 24-hour clock can be provided (month, day, hour, and minutes).

house phone—allows certain stations to reach the attendant or another station by merely going off-hook.

howler tone—tone used to alert a subscriber when the handset is off-hook.

hum—spurious electrical interference, usually picked up from a conventional alternating current power supply.

hunting—movement of a call as it progresses through a group of lines. Typically, the call will try to be connected on the first line of the group; if that line is busy it will try the second line, and then the third, etc.

hybrid circuit—circuit having four sets of terminals arranged in two pairs designed so that there is high loss between the two sets of terminals of a pair when the terminals

Glossary

of the other pair are suitably terminated. Commonly used to couple four-wire circuits to two-wire circuits.

hyper-density—1-inch magnetic tape with a density of 3,200 characters per inch.

Hz—see *hertz*.

I

I/O—see *input/output*.

I/O bound—refers to programs with a large number of I/O (input/output) operations, which slow the CPU.

IBS—see *international business service*.

ICA—International Communications Association.

ICC—International Control Center.

ICI—see *incoming call identification*.

ICPT—intercept tone. See *vacant number intercept*.

IDD—see *international direct dialing*.

identified outward dialing (IOD)—see *automatic identification of outward dialing (AIOD)*.

IDF—see *intermediate distributing frame*.

IDN—see *integrated digital network*.

IDR—see *international digital route*.

IEC—1) interexchange carrier. 2) International Electrotechnical Commission. Affiliated to the ISO and responsible for international standards in the electrotechnical field. Some work relates to telecommunications, particularly in the area of wires, cables, waveguides, and CATV systems, but concentrating on standards for materials, components, and methods of measurement.

IFRB—see *International Frequency Registration Board*.

IMTS—see *international message telephone service*.

IMV/VS—see *information management system/virtual storage*.

INC—incoming.

incoming call identification (ICI)—allows the attendant to identify visually the type of service, trunk circuit, or trunk group associated with a call.

incoming call indicator—lamp panel and associated audible tone providing an alert signal to personnel for incoming

calls. This unit is usually placed at strategic locations within a customer's premises and is used with console-less systems.

indication of camp-on to station—short bursts of tone periodically transmitted to the busy station to indicate to that station that another call is waiting.

individual transfer, all calls—provides the capability of transferring any call by the station to another internal station. Features of consultation hold and add-on conference are usually included.

induction coil—apparatus for obtaining intermittent high voltage consisting of a primary coil through which the direct current flows, an interrupter, and a secondary coil of a larger number of turns in which the high voltage is induced.

information feedback—see *message feedback*.

information management system/virtual storage (IMS/VS)—a common IBM host operating environment, usually under the MVS operating system, oriented toward batch processing and telecommunications-based transaction processing.

information path—the functional route by which information is transferred.

information systems network—a network of multiple operational-level systems and one management-oriented system (centered around planning, control, and measurement processes). The network retrieves data from databases and synthesizes that data into meaningful information to support the organization.

infrared—pertaining to the frequency range in the electromagnetic spectrum that is higher than radio frequencies but below the range of visible light. See also *electromagnetic spectrum*.

infrared radiation—electromagnetic radiation with wavelengths of between 780 and 10⁵ nanometers.

INMARSAT—International Maritime Satellite Organization. Its work now increasingly covers aeronautical and mobile, as well as maritime, satellite communications.

input/output—a general term for the equipment used to communicate with a computer, commonly called I/O; also the data involved in such communication.

input/output (I/O) channel—a component in a computer system, controlled by the central processing unit, that handles the transfer of data between main storage and peripheral equipment.

Glossary

insertion loss—difference in the amount of power received before and after a device is inserted into a circuit or a call is connected.

Institute of Electrical and Electronics Engineers (IEEE)—a group involved in recommending standards for the computer and communications field.

integrated adapter—provides for the direct connection of an external device and uses neither a control unit nor the standard I/O interface.

integrated circuit—a combination of the interconnected circuit elements inseparably associated on or within a continuous substrate; see also *substrate*.

integrated digital network (IDN)—network employing both digital switches and digital transmission.

integrated services digital network (ISDN)—project underway within the CCITT for the standardization of operating parameters and interfaces for a network that will allow a variety of mixed digital transmission services to be accommodated. Access channels under definition include a basic rate defined by CCITT 2B + D (64K + 64K + 16K bps, or 144K bps) and a primary rate that is DS1 (1.544M bps in the U.S., Japan, and Canada, and 2.048M bps in Europe).

integrated voice/data terminal (IVDT)—devices that feature a keyboard/display and a voice telephone instrument; many contain varying degrees of local processing power, ranging from full PC capacity to directory storage for automatic telephone dialing. They may be designed to work with a specific customer-premises PBX or be PBX-independent.

Intel—Intel Corporation, a major microprocessor and computer systems manufacturer.

intelligent terminal—a terminal with some logical capability; a remote device that is capable of performing some editing or other function upon input or output data; contrast with *dumb terminal*.

Intelsat—International Telecommunications Satellite Consortium, formed in 1964 with the purpose of creating a worldwide communications satellite system.

inter-PBX call transfer—part of a main/satellite configuration. An incoming exchange call to a main PBX or a satellite PBX can be put in a three-way conference mode. In addition, an incoming exchange call to a main PBX can be transferred over a tie trunk to a satellite PBX station and vice versa.

inter-PBX coordinated station numbering—component of main/satellite configurations. Stations at the main and sat-

ellite can dial each other without intervening dial tone. The dialing plan for an inter-PBX call is the same as for an intra-PBX (station-to-station) call.

interactive—1) in communications, describing time-dependent data communications, typically where a user enters data and then awaits a response message from the destination before continuing. 2) conversational.

interchange point—a location where interface signals are transmitted.

intercom—internal communications system that allows calling generally within the same building, but not outside the system. Key systems are frequently provided with intercom lines that allow quick communication between stations on the key system. The PBX has features that can replace or enhance the familiar key system intercom functions on custom key sets.

interconnect company—organization, other than the serving telephone company, that supplies telephone equipment by sale, rental, or leasing.

interexchange channel (IXT)—direct circuit between exchanges.

interface—boundary between two pieces of equipment across which all the signals that pass are carefully defined. The definition includes the connector signal levels, impedance, timing, sequence of operation, and the meaning of signals.

interface EIA standard RS-232C—a standardized method adopted by the Electronic Industries Association to insure physical and signal uniformity of interface between data communication equipment and data processing terminal equipment.

interLATA—between different Local Access and Transport Areas.

intermediate distributing frame (IDF)—frame having distributing blocks on both sides, permitting connection of any telephone number with any line circuit.

international business service—a satellite-based service at up to 8M bps. Services include data, facsimile, digital voice, and video- and audioconferencing.

international digital route—a proposed digital service intended to replace the current analog system. It is designed to allow access to the public switched network at rates up to 45M bps.

international direct dialing (IDD)—cooperative service enabling subscribers to place international calls without operator assistance.

Glossary

International Frequency Registration Board (IFRB)—within the International Telecommunication Union (ITU), the IFRB is responsible for the maintenance of an international list of radio frequency usage and the allocation of new frequencies.

international message telephone service—a public switched voice service.

international number—digits that have to be dialed after the international prefix to call a subscriber in another country; i.e., the country code followed by the subscriber's national number.

International Organization for Standardization (ISO)—an organization established to promote the development of standards to facilitate the international exchange of goods and services, and to develop mutual cooperation in areas of intellectual, scientific, technological, and economic activity.

international record carrier (IRC)—term for group of common carriers that until recently provided data and text service between certain U.S. gateway cities and locations abroad.

international switching center (ISC)—exchange used to switch traffic between different countries over international circuits.

international subscriber dialing (ISD)—see *international direct dialing*.

International Telecommunication Union (ITU)—telecommunications agency of the United Nations, established to provide standardized communications procedures and practices including frequency allocation and radio regulations on a worldwide basis. Parent group of the CCITT.

International Telecommunications Satellite Consortium (Intelsat)—formed in 1964 with the purpose of creating a worldwide communications satellite system.

internet protocol (IP)—provides a common layer over dissimilar networks. It controls communications between two SPU's residing on two or more different networks.

interoffice trunk—direct trunk between local central offices (Class 5 offices), or between Class 2, 3, or 4 offices; also called *intertoll trunk*.

interposition calling—one attendant in a multiposition system can call an attendant at another position for consultation.

interposition transfer—an operator at one console can transfer a call to an operator at another position. Used where certain positions are assigned to handle specific types of calls.

interstate—between different states; over a state line.

intraLATA—within a single Local Access and Transport Area.

intraoffice trunk—trunk connection within the same central office.

intrastate—within a single state's boundaries.

INTUG—see *International Telecommunications Users Group*.

inverter—used to convert a direct current into a higher voltage direct current or an alternating current.

inward restriction—blocks selected extension lines from receiving incoming exchange network calls and CCSA calls, completed via either DID or the attendant. Calls can be given any intercept treatment.

IOD—identified outward dialing (may use operator).

IPM—impulses per minute (interruption rate for call progress tones).

IPSS—international extension of British Telecom's PSS.

IRC—see *international record carrier*.

ISC—see *international switching center*.

ISD—international subscriber dialing. See *international direct dialing*.

ISDN—see *integrated services digital network*.

ISO—see *International Organization for Standardization*.

ISPABX—integrated services PABX, compatible with ISDN standards.

ITU—see *International Telecommunication Union*.

IVDT—see *integrated voice/data terminal*.

IXC—interexchange carrier.

IXSD—international telex subscriber dialing.

IXT—see *interexchange channel*.

Glossary

J

jack—a device used generally for terminating the permanent wiring of a circuit, access to which is obtained by the insertion of a plug.

jitter—slight movement of a transmission signal in time or phase that can introduce errors and loss of synchronization for high-speed synchronous communications. See also *phase jitter*.

Josephson junction—a type of circuit capable of switching at very high speeds when operated at a temperature approaching absolute zero.

jumper—patch cable or wire used to establish a circuit, often temporarily, for testing or diagnostics.

junctor—in crossbar systems, a circuit extending between frames of a switching unit and terminating in a switching device on each frame.

K

K—the symbol for 2, raised to the tenth power (kilo), equal to 1,024.

K bps—kilobits per second.

Ka-Band—portion of the electromagnetic spectrum allotted for satellite transmission; frequencies are approximately in the 20 to 30GHz range.

Kermit—a public-domain file transfer protocol designed at Columbia University that is packet oriented. Kermit uses eight data bits per byte and supports sliding windows and data compression.

key data entry devices—a keyboard-equipped device used to prepare information so that the computer can accept it; includes key punches (card punches), key-to-tape and key-to-disk units.

key illumination: incandescent lamp behind key—circuit status lamp indication behind circuit pushbutton with flashing incoming, steady busy, and “wink” hold visual indications.

key pulsing (KP)—manual method of sending numerical and other signals by the operation of nonlocking pushkeys. Also called *key sending*.

key service unit (KSU)—central operating unit of a key telephone system.

keyboard perforator—perforator provided with a bank of keys, the manual depression of any one of which will cause the code of the corresponding character or function to be punched in a tape.

keyboard send/receive (KSR)—a teleprinter terminal containing a keyboard and printer, but no data storage capability. Messages are transmitted character-by-character as they are keyed, and printed as they are received. Contrast with *automatic send/receive (ASR)*.

keyboard terminal—a station through which data can be entered to a data processing system by means of a typewriter-like keyboard.

keyword—one of the significant and informative words in a title or document that describes the content of that document.

keyword-in-context (KWIC)—an index, which lists available programs in alphabetical order by the most significant word in the title; a KWIC index is prepared by highlighting each keyword of the title in the context of words on either side of it and aligning the keywords of all titles alphabetically in a vertical column.

KHz—see *kilohertz*.

kill—to eliminate or erase; frequently a control (EC) character in a word processing program meaning to drop or purge a line of text or a blank line.

kilobyte—1,024 bytes, or 2^{10} bytes.

kilohertz (KHz)—one thousand cycles per second.

KP—key pulse (signaling unlocking signal). See *key pulsing*.

KSR—see *keyboard send/receive*.

KSU—see *key service unit*.

KTU—key telephone unit. See *key service unit*.

ku band—portion of the electromagnetic spectrum being used increasingly for satellite communications. Frequencies approximately in the 12GHz to 14GHz range.

KW—kilowatt.

KWH—kilowatt hour.

L

L-Band—portion of the electromagnetic spectrum commonly used in satellite and microwave applications, with frequencies approximately in the 1GHz region.

label—set of symbols used to identify or describe an item, record, message, or file. Occasionally it can be the same as the address in storage.

LAMA—see *local automatic message accounting*.

Glossary

LAN—see *local area network*.

language—a set of rules and conventions used to convey information.

large-scale computer—a system with the highest operating characteristics; it provides complex and powerful programmable logic to attack problems that require high amounts of computing power.

large scale integration (LSI)—method of fabricating electronic chips permitting a large number of circuits on a single chip. Also called *very large scale integration (VLSI)*.

laser—acronym for “light amplification by stimulated emission of radiation.” Lasers convert electrical energy into radiant energy in the visible or infrared parts of the spectrum, emitting light with a small spectral bandwidth. For this reason, they are widely used in fiber optic communications, particularly as sources for long-haul links.

laser disk—an analog or digital storage medium written and read by laser; see also *video disk*.

last number redial—a PBX/key system electronic telephone feature in which the system or set stores the last number dialed and automatically redials it when the user presses a designated button.

last trunk busy (LTB)—condition in which the last-choice trunk of a given group is busy.

LATA—see *local access and transport area*.

latency—the time interval between when a network station seeks access to a transmission channel and access is granted or received.

layer—in the OSI reference model, referring to a collection of related network-processing functions that comprise one level of a hierarchy of functions.

LCD—see *liquid crystal display*.

LCR—see *least cost routing, automatic route selection*.

LCU—see *line control unit*.

LD—see *long-distance*.

LDN—see *listed directory number*.

leading zeros—zeroes that have no significance in the value of an arithmetic integer; all zeroes to the left of the first significant integer digit of a number.

leaky PBX—a PBX that allows incoming private-line (tie line) traffic to access the local exchange.

learning curve—a planning technique calculation based on the premise that workers will be able to produce a product more quickly after they get used to making it.

lease—a contract whereby one party, known as the lessor, grants to another party, known as the lessee, the rights to use the property owned by the lessor. This property may be land, buildings, or equipment. The lease agreement describes the rights of the owner (lessor) and renter (lessee) and the terms of payment and the tenure of the lease.

leased line—a communication channel contracted for exclusive use from a common carrier, frequently referred to as a private line.

least cost routing (LCR)—see *automatic route selection*.

least significant character—the extreme right hand character in a group of characters; contrast to *most significant character*.

least squares method—a method of smoothing or curve fitting which selects the fitted curve so as to minimize the sum of squares of deviations from the given points.

LEC—local exchange carrier; see *local exchange*.

LED—see *light-emitting diode*.

letter quality printer—a device that generates output that is suitable for high-quality business correspondence; the term implies that output quality matches that of a standard office typewriter.

letters shift—physical shift in a teletypewriter, specifically telex, which enables the printing of alphabetic characters. Also the name of the character that causes this shift.

level—1) in data management structures or communication protocols, the degree of subordination in a hierarchy. 2) measurement of signal power at specific point in a circuit.

LF—see *low frequency*.

library—a collection of information. In electronic data processing, a program library is a collection of available computer programs and routines. The libraries used by an organization are known as its data bank.

library routine—a proven routine that is stored with other routines.

LIFO (last-in-first-out)—the procedure in data processing in which the newest entry in a queue or file is the first to be removed.

light-emitting diode (LED)—semiconductor junction diode that emits radiant energy and is used as a light source

Glossary

for fiber optic communications, particularly for short-haul links. Also used in alphanumeric displays in electronic telephones and calculators.

light pen—a tool for CRT terminal operators that causes the computer to change or modify the display on the cathode-ray tube. The pen's response to light from the display is transmitted to the computer which, in turn, relates the computer action to the section of the image being displayed. In this way, the operator can delete or add text, maintain tighter control over the program, and choose alternative courses of action.

lightwave communications—term sometimes used in place of "optical" communications to avoid confusion with visual information and image transmission such as facsimile or television; see also *fiber optics*.

limited-distance modem—a device that translates digital signals into analog signals (and vice versa) for transfers over limited distances; some operate at higher speeds than modems that are designed for use over analog telephone facilities.

limiting operation—see *bottleneck*.

line—1) a communications path between two or more points, including a satellite or microwave channel. See *channel*. 2) in data communications, a circuit connecting 2 or more devices. 3) transmission path from a nonswitching subscriber terminal to a switching system. 4) in management structures, an authority relationship in an organization where one person (manager) has responsibility for the activities of another person (subordinate).

line balancing—the assignment of tasks to multiple workstations so as to minimize the number of workstations and the total amount of unassigned time at all stations.

line conditioning—see *conditioning*.

line control unit (LCU)—in data communications, a special-purpose processor that controls input/output of communication lines not directly accessed by the computer.

line discipline—see *control procedure*.

line driver—communications transmitter/receiver which is used to extend the transmission distance between terminals and computers that are directly connected; acts as an interface between logic-circuits and a 2-wire transmission line.

line finder—switch that finds a calling line among a group of lines and connects it to another device. Used typically in step-by-step (S × S) switching.

line hit—electrical interference causing the introduction of undesirable signals on a circuit.

line loading—the process of installing loading coils in series with each conductor on a transmission line. Usually 88 millihenry coils installed at 6,000 foot intervals.

line load control—equipment in a telephone system that provides a means by which essential paths may be assured continuity of service under over-loaded conditions, generally accomplished by temporarily denying originating service to some or all of the nonessential lines.

line lockout with warning—telecommunications system feature which provides 10 seconds of intercept tone and then holds the line out of service when the station line remains off-hook for longer than a predetermined time.

line number—the identification number sequence of an instruction or statement in a sequential program.

line of sight—characteristic of some open-air transmission technologies where the area between a transmitter and a receiver must be clear and unobstructed; said of microwave, infrared, and open-air laser-type transmission; a clear, open-air, direct transmission path free of obstructions such as buildings but in some cases impeded by adverse weather or environmental conditions.

line preference—telecom, system feature in which the user selects a line by pressing the button associated with that line.

line printer—a computer output device that generates an entire line of characters at a time; this principle provides high printing speed.

line processing—processing of transactions as they occur, with no preliminary editing or sorting of them before they enter the system. Contrast with *batch processing*.

line speed—maximum data rate that can be reliably transmitted over a line.

line switching—see *circuit switching*.

linear programming—refers to mathematical techniques for solving a general class of optimization problems through selection of the least costly or most profitable solution.

link—1) physical circuit between two points. 2) conceptual (or virtual) circuit between two users of a packet switched (or other) network that allows them to communicate, even when different physical paths are used.

link layer—the logical entity in the OSI model concerned with transmission of data between adjacent network nodes; it is the second layer processing in the OSI model, between the physical and the network layers; see also *OSI*.

Glossary

link redundancy level—the ratio of actual number of paths to the minimum number of paths required to connect all nodes of a network.

linkage editor—an operating system program that prepares the output of language translators for execution. It combines separately produced modules; resolves cross-references among them; replaces, deletes, and adds control sections; and produces an executable load module.

liquid crystal display (LCD)—a graphic display on a terminal screen using an electroluminescent technology to form symbols or shapes.

LISP (LISP processing)—an interpretive language developed for manipulation of symbolic strings and recursive data; while the language has been developed to aid in the handling of symbolic lists, it can be and has been used successfully in the manipulation of mathematical and arithmetic logic.

list—a data structure in which each item of data can contain pointers to other items; any data structure can be represented in this way, which allows the structure to be independent of the storage of the items.

listed directory number (LDN)—an organization's main telephone number as listed in the local telephone directory. In a PBX system, calls to this number are directed to the attendant who can complete these calls to station lines.

listing—a printout, generally prepared by a language translator, that lists the source language statements of a program.

LLC—see *logical link control*.

load—in computer operations, the amount of scheduled work, generally expressed in terms of hours of work. In programming, to feed data or programs into the computer.

load balancing for station/trunk lines—traffic balancing may be required when system traffic limits are approached. This feature provides the capability, during installation or on an in-service system to change specific station and trunk terminations on the PBX system switching network with minimum installer effort and without requiring number changes for the purpose of balancing the traffic load on the switch network.

load center—a group of workstations that can be considered together for purposes of scheduling.

load sharing—the distribution of tasks among a number of network computers.

loading—adding inductance (load coils) to a transmission line to minimize amplitude distortion.

loading coil—induction device employed in local loops, generally those exceeding 5,500 meters in length, that compensates for the wire capacitance and serves to boost voice grade frequencies. They are often removed for new generation, high-speed, local-loop data services as they can distort data signals at higher frequencies than those used for voice.

local access and transport area (LATA)—geographic regions within the U.S. that define areas within which the Bell Operating Companies (BOCs) can offer exchange and exchange access services (local calling, private lines, etc.).

local area network (LAN)—a system for linking terminals, programs, storage, and graphic devices at multiple workstations over relatively small geographic areas.

local automatic message accounting (LAMA)—combination of automatic message accounting equipment, automatic number accounting equipment, and automatic number identification equipment in the same office. In such a system, a subscriber-dialed call can be automatically processed without operator assistance.

local call—any call for a destination within the local service area of the calling station.

local call billing—telecom system feature used in hotel applications which computes the charges for local calls placed by guests based on total message units and, optionally, service charges stored for each guestroom telephone via the station message register service feature and the hotel local call billing rate parameter.

local exchange—switching center in which subscribers' lines terminate. The exchange has access to other exchanges and to national trunk networks. Also called *local central office, end office*.

local loop—a line connecting a customer's telephone equipment with the local telephone company exchange.

local service area—area within which the telephone operating company uses local rates for calling charge.

local systems environment—those resources which exist in a real, open system but which are outside the scope of the OSI environment.

local trunk—trunks between local exchanges.

lockout—denies the attendant the ability to reenter an incoming exchange connection directly terminated or held on his/her position, unless specifically recalled by the station user.

logic circuit—an electronic path that performs information processing; generally encoded on a chip, it is composed of logic gates which are the boolean logic building blocks. See *logic gate*.

Glossary

logic design—the specification of the working relations between the parts of a system in terms of symbolic logic.

logic gate—a combination of transistors that detects the presence or absence of electrical pulses, which in turn represent binary digits (1s and 0s).

logical data independence—the capacity to change the overall logical structure of the data without changing the application programs.

logical link control (LLC)—a protocol developed by the IEEE 802, common to all of its local network standards, for data link-level transmission control. The upper sublayer of the IEEE layer-2 (OSI) protocol that complements the MAC protocol (IEEE 802.2).

logical record—a collection of items independent of their physical environment; portions of the same logical record may be located in different physical records.

logical terminal—a device addressable by its logical function rather than its physical address; translation from logical to physical addresses is achieved by a common routine using a table.

log-in—see *logon*.

logo—an interactive programming language developed by Seymour Papert at MIT, primarily for students using an online terminal or PC.

log-off—the procedure by which a user ends a terminal session.

logon—the procedure by which a user begins a terminal session.

logout—see *log-off*.

long-distance—any telephone call to a destination outside the local service area of the calling station. Also called *toll call*.

long-haul—long distance, describing (primarily) telephone circuits that cross out of the local exchange.

longitudinal balance—measure of the electrical balance between the two conductors (tip and ring) of a telephone circuit. Specifically, the difference between the tip-to-ground and ring-to-ground AC signal voltages, expressed in decibels.

longitudinal redundancy check—a data communications error trapping technique in which a character is accumulated at both the sending and receiving stations during the transmission and is compared for an equal condition, which indicates a good transmission of the previous block.

loop—1) local circuit between an exchange and a subscriber telephone station. Also called *subscriber loop* and *local line*. 2) in programming, a sequence of computer instructions that repeats itself until a predetermined count or other test is satisfied.

loop checking—see *message feedback*.

loop circuit—generally refers to the circuit connecting the subscriber's equipment with local switching. Also called *metallic circuit* and *local loop*.

loop signaling systems—any of three methods of transmitting signaling information over the metallic loop formed by the trunk conductors and the terminating equipment bridges. Transmission of the loop signals can be accomplished by 1) opening and closing the DC path around the loop, 2) reversing the voltage polarity, or 3) varying the value of the equipment resistance.

loop start—most commonly used method of signaling an off-hook condition between an analog phone set and a switch, where picking up the receiver closes a wire loop, allowing DC current to flow, which is detected by a PBX or local exchange and interpreted as a request for service.

loopback—a diagnostic procedure used for transmission devices; a test message is sent to a device being tested, which then sends the message back to the originator for comparison with the original transmission. Loopback testing may be performed within a locally attached device or conducted remotely over a communications circuit. See *analog loopback*, *digital loopback*.

loss (transmission)—decrease in energy of signal power in transmission along a circuit due to the resistance or impedance of the circuit or equipment.

loudspeaker paging access—allows the PBX/key system attendant and station users to have access to paging equipment for the purpose of voice paging. All voice paging facilities make use of the telephone transmitter as the microphone. Available to station users on a class-of-service basis, this function allows direct connection to the paging system by dialing a unique number for each zone or for all zones simultaneously.

low frequency (LF)—generally indicates frequencies between 30 KHz and 300 KHz.

low-level language—a programming language in which instructions have a 1-to-1 relationship with machine code; see also *computer language*.

lpm—lines per minute, a reference to printer speeds.

LRC—see *longitudinal redundancy check*.

LSI—see *large scale integration*.

Glossary

LTB—see *last trunk busy*.

LTE—local telephone exchange; see *local exchange*.

M

M—mega; 1,048,576 bytes (1,024K squared) or roughly 1 million.

Mbps—megabits per second.

machine language—a set of commands written using 0s and 1s so the computer can understand and process information.

machine-oriented code—instructions written using machine language; see *machine language*.

machine sensible—programs and information that can be read and understood by a computer without translation.

machine utilization—the amount of time a system is running as opposed to being idle.

macro—a subroutine that retrieves frequently used phrases and formats with a few keystrokes.

magnetic bubble memory—a nonvolatile storage device that records information in charged areas on a semiconductor plane.

magnetic core—material used to store data in main memory.

magnetic disk—a flat circular plate where data can be stored; the information is accessible to reading and writing heads on an arm which can be moved to a desired storage area as the plate rotates.

magnetic storage—devices where information is stored by electrically charging metal or coated plastic tape.

magnetic tape—flexible plastic material magnetically coated on one side to store information.

mailgram—hard-copy message communications service jointly provided by Western Union and the U.S. Postal Service. Messages received at Western Union offices or originated at Telex/TWX stations are sent over the WU network to a post office near the addressee for next day postal delivery.

main distribution frame (MDF)—a wiring arrangement that connects outside lines on one side and internal lines from exchange equipment on the other.

main memory—the primary storage location of a computer; all information stored there is automatically accessible to the computer; contrast with *external storage*. Also called *main storage*.

main program—a group of instructions necessary for the computer to execute specific instructions.

main/satellite service—a PBX networking feature in which multilocation PBX customers can concentrate their attendant positions at one location, referred to as the main. Other unattended locations, equipped with dial switching equipment, are referred to as satellites. Only one listed directory number (LDN) is provided per complex, and all attendant-handled calls are switched through the main PBX over tie trunks to and from satellite PBX locations.

main station—subscriber's instrument (e.g., telephone or terminal) connected to a local loop, used for originating calls and answering incoming calls from the exchange.

mainframe—see *central processing unit*.

maintenance time—the interval required for hardware correction and enhancement; contrast with *available time*.

makeup time—the interval used for reruns due to malfunctions or mistakes.

malfunction—an operation failure in a machine's hardware.

managed object—a data processing or data communications resource that may be managed through the use of an OSI Management protocol. The resource itself need not be an OSI resource. A managed object may be a physical item of equipment, a software component, some abstract collection of information, or any combination of all three.

management information—information, associated with a managed object that is required to control and maintain that object.

management information base (MIB)—a conceptual composite of information about all managed objects in an open system.

management information system—the collection of computers and data that is used to track and run a company or organization.

manual exchange—exchange in which connections are made by an operator.

manual exclusion—method by which a PBX station user, by entering a special code, can block all other stations on that line from entering the call, assuring privacy on the line.

manual hold—method of placing a line circuit on "hold" by activating a nonlocking common hold button.

manual input—information entered by hand by an operator or programmer to modify, continue, or resume processing of a computer program.

Glossary

manual originating line service—provides station lines that require the attendant to complete all outgoing calls. All nonattendant-handled call attempts are intercepted. This arrangement can be used for lobby phones or emergency telephones to minimize system abuse.

manual PBXs—PBX systems that are not automatic and that require all calls to be connected through the attendant position, including station-to-station calls.

manual signaling—depressing a specific button on a telephone to send an audible signal to a predetermined station.

manual terminating line service—provides extension lines that require all terminating calls to be completed by the attendant. Intercepts all call attempts not handled by the attendant. This feature can be activated, for example, on patient phones in a hospital to prevent disturbance.

manufacturing automation protocol (MAP)—application-specific protocol based on Open Systems Interconnection (OSI) standards. It is designed to allow communication among competing vendors' digital products in the factory.

MAP—see *manufacturing automation protocol*.

mapping—in network operations, the logical association of one set of values, such as addresses on one network, with quantities or values of another set, such as devices on a second network (e.g., name-address mapping, internetwork-route mapping).

MARECS—European maritime satellite system

MARISAT—maritime satellite service.

mark—the signal (communications channel state) corresponding to a binary 1. The marking condition exists when current flows (current-loop channel) or when the voltage is more negative than -3 volts (EIA RS-232-C channel).

marker—wired-logic control circuit that, among other functions, tests, selects, and establishes paths through a switching state(s) in response to external signals.

mass storage—external disk or tape devices capable of holding a very large capacity of information.

master control program—the set of instructions used to initialize and control the computer's memory.

master file—the definitive collection of related records stored in a computer system.

master number hunting—an incoming call routing pattern in which the specific number is assigned as the first number in the hunt group, rather than the more traditional procedure of using the first station within the group.

master station—unit having control of all other terminals on a multipoint circuit for purposes of polling and/or selection.

master terminal—a reserved workstation used by the system manager that can initiate conversations for network management tasks.

mathematical programming—a procedure for locating the maximum or minimum of a function.

matrix—1) an arrangement of elements (numbers, characters, dots, diodes, wires, etc.) in perpendicular rows. 2) in switch technology, that portion of the switch architecture where input leads and output leads meet, any pair of which may be connected to establish a through circuit. Also called a *switching matrix*.

matrix printer—a device that forms characters by outputting a series of dots on paper.

MDF—see *main distribution frame*.

MDNS—managed data network services. A pan-European public VANs operator initiated by the Commercial Action Committee of CEPT.

mean—the arithmetical average value of a group of values; see also *median, mode*.

mean-time-between-failure (MTBF)—average length of time for which the system, or a component of the system, works without fault.

mean-time-to-repair (MTTR)—the average time required to perform corrective maintenance on a failed device.

measured rate—message rate structure which includes payment for a specified number of calls within a defined area, plus a charge for additional calls.

media—the vehicles that store or transmit information, classified as source, input, and output.

median—the middle value in a group.

medium—1) the material on which data is recorded; for example, magnetic tape, diskette. 2) any material substance that is, or can be, used for the propagation of signals, usually in the form of modulated radio, light, or acoustic waves, from one point to another, such as optical fiber, cable, wire, dielectric slab, water, air, or free space (ISO).

medium frequency (MF)—frequencies in the range between 300 KHz and 3 MHz.

medium scale integration—generally less than 100 circuits built onto a single chip; used frequently in third generation systems.

Glossary

meet-me conference—conference circuit on a PBX given a single-digit access code. All stations dialing that code at a predetermined time (or upon direction by the operator) are connected in conference.

megabyte—approximately 1,000,000 bytes.

megahertz (MHz)—a unit of frequency equal to 1 million cycles per second.

memory—area of a computer system that accepts, holds, and provides access to information.

memory address—a coded designation for a location of stored information.

memory dump—a listing of the contents in a storage device.

memory location—a position in a computer's storage area.

memory management—allocating storage space to maximize the amount of information that can be held.

memory map—a graphic representation of all storage locations that a computer can designate.

memory protection—see *storage protection*.

menu—a list of commands and available options in a program that are available to a user.

menu-driven—set of instructions using a list of commands and available options; the user only has to select the desired option; compare to *command-driven*.

merge—an operation combining two or more files of information into one in a predetermined sequence.

message—sequence of characters used to convey information or data. In data communications, messages are usually in agreed format with a heading, which establishes the destination of the message, and text, which consists of the data being sent.

message feedback—method of checking the accuracy of data transmission in which the received data are returned to the sending end for comparison with the original data, which are stored there for this purpose. Also called *information feedback* and *loop checking*.

message format—rules for the placement of such portions of a message as message heading, address, text, end-of-message indication, and error-detecting bits.

message numbering—identification of each message within a communications system by the assignment of a sequential number.

message registration—a hotel telecommunications system feature provides an electronic or mechanical readout of outgoing local and long-distance calls from each guest station. This information can be displayed at a hotel console or on mechanical counters.

message relay—a feature that allows a PBX extension user to record and store a message for transmission to a given extension by a given time, after which it is relayed to the attendant.

message switching—a technique that transfers messages between points not directly connected. The switching facility receives messages, stores them in queues for each destination point, and retransmits them when a facility becomes available. Synonymous with *store-and-forward*.

message unit (MU)—unit of measure for charging local calls that details the length of call, distance called, and time of day.

message waiting—feature that allows an attendant or voice messaging system to light an indicator lamp on a user's telephone to show that a message is waiting.

metal oxide semiconductor—a common material and method for making integrated circuits.

MF—1) medium frequency. 2) multifrequency. See *dual tone multifrequency signaling (DTMF)*.

MHz—see *megahertz*.

MIB—see *management information base*.

MICR (magnetic ink character recognition)—special encoded symbols that can be read and converted to digital information for a computer; generally used on checks and deposit slips.

MICR reader/sorter—a device that translates symbols to digital information that can be understood by a computer; also sorts the original document based on the unique symbols.

micro floppy—see *microdiskette*.

Microcom Networking Protocol (MNP)—an error-correction method based on the OSI reference model that retransmits lost or corrupt data in packets of information. The protocol was originally developed in 1982 by Gregory Pearson of Microcom and now consists of nine separate classes. MNP supports interactive and file transfer applications over dial-up lines at transfer rates up to 38.4K bps.

microcomputer—a small-scale programmable machine that processes information; it generally has a single chip as its central unit and includes storage and input/output facilities in the basic unit.

Glossary

microcomputer kit—small-scale programmable machines marketed in unassembled form.

microdiskette—magnetic storage medium enclosed in a plastic case that is less than 4 inches in size; the most popular size is 3.5 inches.

microfiche—a rectangular sheet of transparent film that contains multiple rows of greatly reduced page images of reports, catalogs, rate books, etc.

microfilm—a small roll of photographic film which can hold several thousand document pages which, when projected onto a screen, produces a legible copy of the item or form photographed.

microprocessor—the central unit of a microcomputer that contains the logical elements for manipulating and performing arithmetical and logical operations on information.

microprogramming—a method of operation of the CPU in which each complete instruction starts the execution of a sequence of instructions, called microinstructions, which are at a more elementary level.

microsecond—one-millionth of a second.

microwave—1) portion of the electromagnetic spectrum above about 760 MHz. 2) describing high-frequency transmission signals and equipment that employ microwave frequencies, including line-of-sight open-air microwave transmission and, increasingly, satellite communications.

micro-Winchester disk—a small, sealed disk under 4 inches in diameter; see also *Winchester disk*.

millisecond—one-thousandth of a second.

min—minimum.

mini floppy—see *minidiskette*.

mini-Winchester disk—a 5.25-inch sealed disk; see also *Winchester disk*.

minicomputer—small- to medium-scale programmable machine that processes information; generally the midrange between microcomputers and mainframes, some can support up to several hundred user terminals simultaneously.

minidiskette—5.25-inch magnetic storage medium.

MIS—see *management information system*.

mixed station dialing—indicates the ability of the switching system to accommodate both rotary dial and pushbutton dial stations.

mnemonic—a symbolic name given to programs, data, and instructions; for example, ART for arithmetic expression, MPY for multiply, and VNDMSTFLE for Vendor Master File.

mobile earth station—radio transmitter and/or receiver situated on a ship, vehicle, or aircraft and used for satellite communications.

mobile satellite communications—satellite-based services to serving ships, motor vehicles, and aircraft.

mobile telephone exchange (MTE)—exchange providing service to mobile telephone subscribers.

mobile telephone service (MTS)—telephone service provided between mobile stations and the public switched telephone network; radio transmission provides the equivalent of a local loop. See also *cellular radio*.

mode—1) the most common or frequent value in a group of values. 2) one of many energy propagation patterns through a waveguide.

mode of operation—the method of processing data to meet a company's need.

model—an approximate mathematical representation that simulates the behavior of a process device or concept so that an increased understanding of the system is attained.

modem—a contraction of *modulate* and *demodulate*; a conversion device installed in pairs at each end of an analog communications line. The modem at the transmitting end modulates digital signals received locally from a computer or terminal; the modem at the receiving end demodulates the incoming signal, converting it back to its original (i.e., digital) format, and passes it to the destination business machine.

modem sharing unit (MSU)—a device that permits two or more terminals to share a single modem.

modular—a design technique that permits a device or system to be assembled from interchangeable components; permits the system or device to be expanded or modified simply by adding another module.

modulation—the process of converting voice or data signal for transmission over a network.

modulation, amplitude (AM)—form of modulation in which the amplitude of the carrier is varied in accordance with the instantaneous value of the modulating signal.

modulation, frequency (FM)—form of modulation in which the instantaneous frequency of a sine wave carrier is

Glossary

caused to depart from the carrier frequency by an amount proportional to the instantaneous value of the modulating signal.

modulation, pulse amplitude (PAM)—form of modulation in which the amplitude of the pulse carrier is varied in accordance with successive samples of the modulating signal.

modulation, pulse code (PCM)—digital transmission technique that involves sampling of an analog information signal at regular time intervals and coding the measured amplitude value into a series of binary values, which are transmitted by modulation of a pulsed, or intermittent, carrier. A common method of speech digitizing using 8-bit code words, or samples, and a sampling rate of 8KHz.

modulation, pulse width (PWM)—process of encoding information based on variations of the duration of carrier pulses. Also called *pulse duration modulation* or *PDM*.

modulator—device that converts a signal (voice or other) into a form that can be transmitted.

module—a hardware or software component that is discrete and identifiable.

module testing—the verification of a discrete and identifiable hardware or software component, generally performed in an isolated environment.

monitor—a software tool used to supervise, control, or verify the operations of a system.

monitor (display)—a device used to display computer-generated information.

monitoring key—feature button that allows an attendant to listen on a circuit without sensibly affecting the transmission quality of that circuit.

morse code—two-condition telegraph code in which characters are represented by groups of dots and dashes.

MOS—see *metal oxide semiconductor*.

most significant character—the extreme left-hand symbol in a group of symbols; contrast to *least significant character*.

motherboard—the central card of a computer; it features female connectors that accept other printed circuit cards.

mouse—a handheld device, separate from a keyboard, used to control the position indicator on a display screen; as the device is rolled along a tabletop, its relative position approximates the position of the indicator.

moving average—the mean of the most recent observations; as each new observation is added, the oldest one is dropped.

MS-DOS—Microsoft Disk Operating System; a master control program for 16-bit systems; see also *Master Control Program*.

msec—millisecond.

MSG—see *message*.

MSU—see *modem sharing unit*.

MTBF—see *mean-time-between-failure*.

MTE—see *mobile telephone exchange*.

MTS—see *mobile telephone service*.

MTTR—see *mean-time-to-repair*.

MTX—see *mobile telephone exchange*.

MU—see *message unit*.

mu-law encoding—encoding according to CCITT recommendation G.711, used with 24-channel PCM systems in the U.S. It is similar to a-law encoding, but the two differ in the size of the quantizing intervals.

multiaccess—the ability for several users to communicate with a computer at the same time.

multibus—Intel's central path (channel) for transmitting electrical signals and data; it was developed for use in 8- and 16-bit computer systems.

multidrop—a communications arrangement where multiple devices share a common transmission channel, though only one may transmit at a time. See also *multipoint*.

multidrop line—see *multipoint line*.

multileaving—the transmission of a variable number of datastreams between user devices and a computer, usually via bisync facilities and using bisync protocols.

multimode fiber—a fiber supporting propagation of multiple modes. See *modes*.

multiple—system of wiring so arranged that a circuit, a line, or a group of lines are accessible at a number of points. Also called *multipoint*.

multiple console operation—PBX supporting more than one attendant's position to handle heavy traffic. Call traffic is distributed evenly among consoles in use.

Glossary

multiple customer group operation—PBX that can be shared by several different companies, each having separate consoles and trunks. Stations are assigned to one company or the other and are then able to reach only that company's trunks and attendants.

multiple listed directory number service—allows more than one listed directory number to be associated with a single PBX installation. Each listed number can be assigned a unique incoming call identification.

multiple trunk groups—indicates that the switching system is capable of being equipped for more than one group of trunk circuits.

multiplex—to interleave or simultaneously transmit two or more messages on a single channel.

multiplex hierarchy—12 channels = 1 group; 5 groups (60 channels) = 1 supergroup; 10 supergroups (600 channels) = 1 mastergroup (U.S. standard); 5 supergroups (300 channels) = 1 mastergroup (CCITT standard); 6 U.S. mastergroups = 1 jumbo group.

multiplexer (mux)—a device that enables more than one signal to be sent simultaneously over one physical channel. It combines inputs from two or more terminals, computer ports, or other multiplexers, and transmits the combined datastream over a single high-speed channel. At the receiving end, the high-speed channel is demultiplexed, either by another multiplexer or by software.

multiplexing—division of a transmission facility into two or more channels either by splitting the frequency band transmitted by the channel into narrower bands, each of which is used to constitute a distinct channel (frequency-division multiplex), or by allotting this common channel to several different information channels, one at a time (time-division multiplexing).

multiport—pertaining or referring to a communications line to which three or more stations are connected. It implies that the line physically extends from one station to another until all are connected. Contrast with *point-to-point*.

multiport line—a communications line to which several stations are connected, though only one can transmit at a time. Usually requires a polling mechanism under the control of a master station to ensure that only one device transmits data at a time. Also called *multidrop*. Contrast with *point-to-point*.

multiprocessing—simultaneous application of more than one processor in a multi-CPU computer system to the execution of a single user job, which is possible only if the job can be effectively defined in terms of a number of independently executable components.

multiprocessing system—a computer capable of performing more than one task at a time.

multiprocessor—a computer system with two or more central computers under common control.

multiprogramming—a technique allowing several programs to run on one computer system at the same time.

multitasking—two or more program segments running in a computer at the same time.

multithread operation—the processing of several distinct communications messages simultaneously.

multithreading—concurrent processing of more than one message (or similar service request) by an application program.

multiuser—a computer that can support several workstations operating simultaneously.

music on hold, system—availability of audio source input for system wide distribution to all "held call" conditions within the system, for attendant and station use.

music on hold, trunk—availability of audio source input for all "held" conditions placed on trunk circuits in the system.

mux—see *multiplexer*.

N

nailed-up connection—slang term for permanent, dedicated path through a switch; often used for lengthy, regular data transmissions going through a PBX.

nak—see *negative acknowledgement*.

nanosecond—one-billionth of a second.

narrowband channels—subvoice grade pathways characterized by a speed range of 100 to 200 bps.

NARUC—National Association of Regulatory Utility Commissioners.

NATA—North American Telecommunications Association.

national number—digits that have to be dialed following the trunk prefix to call a subscriber in the same country but outside the local numbering area. These digits uniquely identify a station in an area identified by a country code.

National Television Systems Committee (NTSC)—television broadcasting system using 525 picture lines and a 60Hz field frequency. Developed by the Committee, and

Glossary

used primarily in the U.S., Canada, Mexico, and Japan. See also *phase alternate line (PAL)* and *Sequential Couleur à Mémoire (SECAM)*.

natural language—a set of commands that use standard English conventions.

NBS—see *NIST*.

near-end crosstalk (NEXT)—unwanted energy transferred from one circuit usually to an adjoining circuit. It occurs at the end of the transmission link where the signal source is located, with the absorbed energy usually propagated in the opposite direction of the absorbing channel's normal current flow. Usually caused by high-frequency or unbalanced signals and insufficient shielding.

negative acknowledge character—a transmission control symbol that indicates a block of information was received incorrectly.

negative acknowledgement (NAK)—an indication that a previous transmission was in error and the receiver is ready to accept retransmission.

neper—basic unit of a logarithmic scale used for the expression of ratios of voltages, currents, and similar quantities.

NETS—Normes Europeennes de Telecommunications. Pan-European communications equipment type approval standards agreed to by the European Telecommunications Standards Institute (ETSI).

network—1) a series of points connected by communications channels. 2) switched telephone network is the network of telephone lines normally used for dialed telephone calls. 3) private network is a network of communications channels confined to the use of one customer.

network architecture—the philosophy and organizational concept for enabling communications between data processing equipment at multiple locations. The network architecture specifies the processors and terminals, and defines the protocols and software that must be used to accomplish accurate data communications.

network control program—an interface routine that coordinates the communications pathway and user programs on the host computer; see also *interface*.

network database—an organization method that allows information relationships to be expressed.

network file system (NFS)—an industry standard for remote file access across a common network; it allows workstations to share file systems in a multivendor network of machines and operating systems. Includes RFS, RPC, XDR, and YP.

network file transfer (NFT)—copies files between any two nodes on a network, either interactively or programmatically. Allows user to 1) copy remote files, 2) translate file attributes, and 3) access remote accounts.

network interprocess communications (NetIPC)—permits inter-CPU program sharing and allows programs running on different systems to exchange data through a set of programmatic calls. This peer-to-peer service is important for developing distributed applications.

network inward/outward dialing (NIOD)—ability to provide dialing both ways between a toll network and a local network.

network layer—in the OSI model, the logical network entity that services the transport layer. It is responsible for ensuring that data passed to it from the transport layer is routed and delivered through the network.

network management center (NMC)—center used for control of a network. May provide traffic analysis, call detail recording, configuration control, fault detection and diagnosis, and maintenance.

network redundancy—a communications pathway that has additional links to connect all nodes.

network security—measures taken to protect a communications pathway from unauthorized access to and accidental or willful interference of regular operations.

network terminating unit (NTU)—the part of the network equipment that connects directly to the data terminal equipment.

network topology—describes the physical and logical relationship of nodes in a network; the schematic arrangement of the links and nodes of a network, typically either a star, ring, tree, or bus topology, or some hybrid combination thereof.

network virtual terminal—a communications concept wherein a variety of data terminal equipment (DTEs), with different data rates, protocols, codes, and formats, are accommodated in the same network. This is done as a result of network processing, where each device's data is converted into a network standard format, then converted into the format of the receiving device at the destination end.

networking—the connection of geographically separated computers using transmission lines.

NEXT—see *near-end crosstalk*.

night audit—provides automatic printout of message registration data for all guestrooms by key operation at front desk console.

Glossary

night console position—provides an alternate attendant position that can be used at night in lieu of the regular console.

night service automatic switching—should the attendant neglect to place the console in the night answering mode, after a certain period of timed ringing from an incoming exchange call, the entire system will automatically engage the night service mode.

night station service—1) fixed service—provides arrangements to route calls normally directed to the attendant to preselected station lines within the PBX system when regular attendant positions are not in use. In addition, calls to specific trunks can be arranged to ring on specific station lines. The receiving station can then transfer the call if necessary. 2) expanded service—routes calls normally directed to the attendant to preselected station lines within the system when it is arranged for night service. Calls to specific exchange trunks can be arranged to route to specific station lines and can be assigned on a flexible basis. Trunk Answer From Any Station capability is provided for calls that are not handled by assigned night stations.

NIOD—see *network inward/outward dialing*.

NIST—National Institute of Standards and Technology. Prior to 1988, was called the National Bureau of Standards.

nmb—number.

NMC—see *network management center*.

NMOS—n-channel metal oxide semiconductor.

NND—national number dialing; see *national number*.

NNX—older form of central office code where N is any digit Room 2 to 9 and X is any digit from 0 to 9.

node—a termination point for two or more communications links. The node can serve as the control location for forwarding data among the elements of a network or multiple networks, as well as performing other networking and, in some cases, local processing functions. A node is usually connected to the backbone network and serves end points and/or other nodes.

noise—unwanted electrical signals, introduced by circuit components or natural disturbances, that tend to degrade the performance of a communications channel. Contrast with *signal*.

nonblocking—describes a switch where a through traffic path always exists for each attached station. Generically, a switch or switching environment designed to never experience a busy condition due to traffic volume.

nonconsecutive hunting—nonconsecutive station numbers can be “searched” by the switching equipment upon dialing the initial number in the hunting group to find connection to the first nonbusy station.

nonvolatile storage—a storage medium whose contents are not lost when the power is removed.

NPA—see *numbering plan area*.

NRZ—non-return to zero (magnetic tape format).

ns—nanosecond (also nsec).

NTSC—National Television Systems Committee. Television broadcasting system using 525 picture lines and a 60Hz field frequency. Developed by the Committee, and used primarily in the U.S., Canada, Mexico, and Japan. See also *PAL* and *SECAN*.

NTU—see *network terminating unit*.

number-unobtainable tone—audible signal received by the caller indicating that the attempted call cannot be completed due to faulty equipment or lines, invalid number dialing, or because access to that number is barred.

numbering plan area (NPA)—geographic subdivision of the territory covered by a national or integrated numbering plan. An NPA is identified by a distinctive area code (NPA code).

numbering zone—one of the nine geographical areas of the world in the world numbering plan.

numeric database—this type of source database typically contains numeric values from original sources and/or data that has been summarized or otherwise statistically manipulated. Most often presented in the form of time series, numeric databases range from simple balance sheet data to complex econometric models. Economic, demographic, and financial data are often depicted in numeric databases.

NXX—a central office code (exchange code) of three digits that designates a particular central office or a given 10,000-line unit of subscriber lines; N is any number from 2 to 9, and X is any number from 0 to 9.

O

object code—see *object program*.

object program—a fully compiled or assembled program that is ready to be loaded into the computer; the result of processing a source program through an assembler or compiler. Also called *object code*.

OCB—see *outgoing calls barred*.

Glossary

OCC—Other Charges or Credits (on phone bill), also Other Common Carrier (MCI, U.S. Sprint, etc.).

octet—8-bit byte. Term used in preference to byte when talking about packet services.

ODA—office document architecture. An OSI standard intended to provide parameters for electronically transmitted documents.

ODD—see *operator distance dialing*.

OEM—see *original equipment manufacturer*.

off-hook—telephone set in use; handset is removed from its cradle.

off-line—1) pertaining to equipment or devices not under direct control of the central processing unit. 2) used to describe terminal equipment that is not connected to a transmission line.

off-premises extension (OPX)—telephone extension located other than where the main switch is.

off-the-shelf—equipment already manufactured and available for delivery from stock.

office classification—see *class of exchange*.

OGT—see *outgoing trunk*.

OHQ—off-hook queue. See *queue*.

on-hook—telephone set not in use; handset resting in cradle.

on-hook dialing—station user can dial a number and listen to the call's progress over the set's speaker, leaving the receiver on-hook until the call goes through and conversation begins.

on-line—1) connected to a computer so that data can pass to or from the computer without human intervention. 2) directly in the line loop. 3) in telegraph usage, transmitting directly onto the line rather than, for example, perforating a tape for later transmission.

online processing—the operation of terminals, disks, and other equipment under complete control of the central processing unit with no need for human intervention.

online services—computer functions offered to end users not owning a host computer; includes timesharing, archival storage, and prepared software programs.

online services company—provides such computer functions as timesharing and prepared software programs to end users connected by terminal to a remote computer.

online software package—program offered by a company with a host computer that manipulates information input by remote end users.

online system—a system with a direct interface between applications programs stored in the computer and terminals for data entry and output.

one-way splitting—when the attendant is in connection with an outside trunk and an internal station, activation of a key allows the attendant to speak privately with the internal station.

one-way trunk—trunk between a switch (PBX) and an exchange, or between exchanges, where traffic originates from only one end.

ONI—see *operator number identification*.

ONP—see *open network provision*.

open network provision (ONP)—a developing pan-European standard for ensuring the provision of the network infrastructure by European telecommunications administrations to users and competitive service providers on terms equal to those for the administrations themselves.

open systems interconnection (OSI)—referring to the reference model, OSI is a logical structure for network operations standardized within the OSI; a seven-layer network architecture being used for the definition of network protocol standards to enable any OSI-compatible computer or device to communicate with any other OSI-compliant computer or device for a meaningful exchange of information.

operating system—software or firmware that controls the internal operation of a computer.

operating time—the time required for dialing the call, waiting for the connection to be established, and coordinating the transaction with the personnel or equipment at the receiving end. Also known as *call setup time*.

operator—person assigned to the operation of a switchboard or comparable equipment. Usually refers to telephone company personnel, but often used interchangeably with the term *attendant*.

operator distance dialing (ODD)—establishment of long-distance calls requiring the intervention of an intermediate operator.

operator number identification (ONI)—at an exchange, a feature that allows the operator to come in long enough to acquire the calling number so that it can be keyed into CAMA equipment.

OPS—off-premises station; see *off-premises extension*.

Glossary

optical character recognition—a light-sensitive scanning process where a device perceives actual character images and converts them into digital code; see also *digital*.

optical disks—storage devices that use laser technology to record data; they feature greater storage capacity than magnetic disks but do not allow data to be over-written.

optical fiber—any filament or fiber, made of dielectric materials, that is used to transmit laser- or LED-generated light signals, usually for digital communications. An optical fiber consists of a core, which carries the signal, and cladding, a substance with a slightly higher refractive index than the core, which surrounds the core and serves to reflect the light signal back into it. Also called *lightguide* or *fiber-optic waveguide*.

OPX—see *off-premises extension*.

OR—originating register (crossbar switching).

original equipment manufacturer (OEM)—maker of equipment that is marketed by another vendor, usually under the name of the reseller. The OEM may only manufacture certain components, or complete computers, which are then often configured with software, and/or other hardware, by the reseller.

originating restriction—station line with this restriction cannot place calls at any time. Calls directed to the station, however, will be completed normally.

OS—see *operating system*.

Osborne 1—the first portable personal computer; see also *personal computer*.

OSI—see *open systems interconnection*.

OSI environment—those resources that enable information processing systems to communicate openly, that is, to conform to the services and protocols of Open Systems Interconnection (OSI).

OSI management—the facilities to control, coordinate, and monitor the resources which enable communications in an OSI environment.

OSITOP—a promotions group for OSI standards.

other common carrier (OCC)—specialized common carriers (SCC), domestic and international record carriers, and domestic satellite carriers engaged in providing services authorized by the Federal Communications Commission.

OTQ—see *outgoing trunk queuing*.

outgoing calls barred (OCB)—prevention of calls to distant addresses; in particular, preventing PBX users from making calls outside the PBX system.

outgoing line restriction—ability of the system to selectively restrict any outside line to an “incoming only” line.

outgoing station restriction—ability of the system to restrict any given station from originating outgoing calls.

outgoing trunk (OGT)—one-way trunk that carries only outgoing traffic.

outgoing trunk queuing (OTQ)—extensions can dial a busy outgoing trunk group, be automatically placed in a queue, and then be called back when a trunk in the group is available. This feature allows more efficient use of expensive special lines as, instead of having to redial the trunk access code until a line is free, the caller can activate OTQ, followed by the trunk access code, and take care of other affairs until a line is free.

output device—external equipment that receives and records information from a storage device or computer; for example, video terminals, printers, and card punches.

outward restriction—station lines within the PBX can be denied the ability to access the exchange network without the assistance of the attendant. Restricted calls are routed to intercept tone.

overflow—excess traffic, on a particular route, that is offered to another (alternate) route.

override—seizure of a circuit even though the circuit is already occupied.

P

PA—public address (loudspeaker system, sometimes used for paging).

PABX—private automatic branch exchange; see *private branch exchange (PBX)*.

packet—group of binary digits, including data and call control signals, that is switched as a composite whole. The data, call control signals, and error control information are arranged in a specified format.

packet assembler/disassembler (PAD)—a protocol conversion device or program that permits end-user devices (e.g., terminals) to access a packet switched network. See also *packet switching*.

packet overhead—a measure of the ratio of the total packet bits occupied by control information to the number of bits of data, usually expressed as a percent.

Glossary

packet switched network—a network designed to carry data in the form of packets. The packet and its format is internal to that network. The external interfaces may handle data in different formats, and conversion is done by an interface computer.

packet switching—a data communications technique where a message is broken down into fixed-length units which are then transmitted to their destination through the fastest route; although all units in a message may not travel the same pathway, the receiving station ascertains that all units are received and in proper sequence before forwarding the complete message to an addressee.

packet-switching network—network designed to carry data in the form of packets. The packet and its format are internal to that network. The external interfaces can handle data in different formats, and conversion is done by an interface computer.

packing density—the amount of storage area available on or in a storage unit; for example, the number of available bytes on a disk.

PAD—see *packet assembler/disassembler*.

paging—see *radio paging*.

paging by zone—by dialing the appropriate access code, any station is able to selectively page groups of predesignated stations or speakers.

paging speakers—inclusion of such speakers within the station instrument. Also includes external units located in larger areas.

paging, total system—upon dialing the appropriate special code, any station can originate a paging announcement through all loudspeakers.

PAL—see *phase alternate line*.

PAM—see *pulse amplitude modulation*.

parallel transmission—the simultaneous transmission of all the bits making up a character or byte, either over separate channels or on different carrier frequencies on the same channel. Contrast with *serial transmission*.

parity—a constant state or equal value. Parity checking is one of the oldest error checking techniques. Character bit patterns are forced into parity (total number of one bits, odd or even) by adding a one or zero bit, as appropriate, as they are transmitted; the parity (odd or even) is then verified upon receipt by the receiving device.

parity bit—a check bit appended to an array of binary digits to make the sum of all the binary digits, including the check bit, always odd or always even.

parity check—addition of noninformation bits to data to make the number of ones in a grouping of bits either always even or always odd. This procedure allows detection of bit groupings that contain single errors. It can be applied to characters, blocks, or any specific bit grouping. Also called *vertical redundancy checking (VRC)*.

party 68—section of FCC rules governing the direct connection of nontelephone company-provided terminal equipment to the telephone network.

party line—subscriber line upon which several subscribers' stations are connected, possibly with selective calling.

party line stations—two-party station service can be expanded to support multiparty service.

patch—1) a temporary electrical connection. 2) to make an improvised modification. To change a software routine in a rough or expedient way.

PATX—see *private automatic telex exchange*.

PAX—see *private automatic exchange*.

PBX—see *private branch exchange*.

PCB—printed circuit board.

PCM—see *pulse code modulation*.

PC-to-host communications products—software or software/hardware products that allow a PC to communicate with a host or mainframe. The most basic PC-to-host product will support terminal emulation capabilities. Today, many of the products also include file-transfer capabilities, allowing PCs to access data that has traditionally been under the control of the mainframe.

PDM—pulse duration modulation; see *pulse width modulation*.

PDN—see *public data network*.

PDX—see *private digital exchange*.

peg count—tally of the number of calls made or received over a specified period.

perforator—instrument for the manual preparation of a perforated tape in which telegraph signals are represented by holes punched in accordance with a predetermined code.

peripheral device or equipment—input or output unit that is not included within the confines of the primary system, e.g., data printer or diskette.

Glossary

permanent virtual circuit—in packet switching, a connection for which a single dedicated path is chosen for a particular transmission. The network is aware of a fixed association between two stations; permanent logical channel numbers are assigned exclusively to the permanent circuit, and devices do not require permission to transmit to each other. A new connection between the same users may be routed along a different path. Contrast with *virtual circuit*.

phantom circuit—third voice circuit, which is superimposed on two 2-wire voice circuits.

phase alternate line (PAL)—color television broadcasting system developed in West Germany and the U.K. that uses 650 picture lines and a 50Hz field frequency. See also *National Television Systems Committee (NTSC)* and *Sequential Couleur à Mémoire (SECAM)*.

phase jitter—a random distortion of signal lengths caused by the rapid fluctuation of the frequency of the transmitted signal. Phase jitter interferes with interpretation of information by changing the timing.

phase modulation (PM)—a way of modifying a sine wave signal to make it carry information. The sine wave, or carrier, has its phase changed in accordance with the information to be transmitted.

phase-shift keying (PSK)—modulation technique for transmitting digital information whereby that information is conveyed as varying phases of a carrier signal.

phosphor—on a CRT, the material which coats the back side of the screen. A dot matrix character, for example, is displayed by illuminating the phosphor of a specified location for each dot defined within its matrix.

photodiode—device consisting essentially of a pin or pin junction diode that converts electromagnetic radiation in the visible or infrared wavelengths into an electric current. Different types are used as detectors in fiber-optic communications, including sensitive avalanche photodiodes.

physical layer—within the OSI model, the lowest level of network processing, below the link layer, that is concerned with the electrical, mechanical, and handshaking procedures over the interface that connects a device to a transmission medium (i.e., RS-232-C).

pilot tone—test frequency of controlled amplitude transmitted over carrier system for monitoring and control purposes.

plant—general term used to describe the physical equipment of a telephone network that provides communications services.

plotters—devices that convert computer output into drawings on paper or on display-type terminals instead of a printed listing.

plug-in stations—station cabling requirement remains constant for all types of station instruments. For ease in station moves and rearrangements, all stations are provided and installed as plug-in instruments.

PM—see *phase modulation*.

PMBX—see *private manual branch exchange*.

PMOS—P-channel metal oxide semiconductor.

POCSAG—a U.K. standard for radio-paging formulated by the Post Office Coding Standards Advisory Group in 1982, now adopted by the CEPT as a pan-European standard.

point of presence (POP)—since divestiture, the physical access location within a LATA of a long-distance and/or inter-LATA common carrier; the point to which the local telephone company terminates subscribers' circuits for long-distance, dial-up, or leased-line communications.

point-to-point line—describing a circuit that connects two points directly, where there are generally no intermediate processing nodes, although there could be switching facilities. Synonymous with *two-point*. See also *multipoint* and *broadcast*.

polar keying—a technique of current loop signaling in which current flow direction establishes the two-level binary code.

polarization—property of an electromagnetic wave characterized by the direction of the electric field.

polling—a method of controlling the sequence of transmission by terminals on a multipoint line by requiring each terminal to wait until the controlling processor requests it to transmit. Contrast with *contention*.

port—a point of access into a communications switch, a computer, a network, or other electronic device; the physical or electrical interface through which one gains access; the interface between a process and a communications or transmission facility.

position—part of a switchboard normally controlled by an operator or attendant.

Postal, Telegraph, and Telephone Organization (PTT)—usually a governmental department that acts as its nation's common carrier.

POTS—plain old telephone service.

Glossary

power failure—PBX users attempting to cope with commercial power failures have a number of alternatives available, ranging from setting up alternative power sources to arranging for the system to fail gradually. See *brownout operation, reserve power, uninterruptible power supply (UPS)*.

power failure transfer—if the PBX is unable to get enough power, this feature provides service to and/or from the exchange network for a limited number of prearranged stations at the customer location. This feature is not available with DID service. Power failure stations usually have an external button that will establish dial tone during outage.

PPS—pulses per second.

prefix—digits that must be dialed to indicate that a call is directed outside the local area.

presentation layer—in the OSI model, that layer of processing that provides services to the application layer, allowing it to interpret the data exchanged, as well as to structure data messages to be transmitted in a specific display and control format.

preset call forwarding—incoming calls are rerouted to a predetermined secondary number.

preventive maintenance—the routine checking of components to keep the system functioning.

primary block—see *primary group*.

primary carrier—long-distance carrier selected by a subscriber as the first-choice provider of long-distance service. Calls placed through the primary carrier require no additional digits, while calls placed with other carriers require that a five-digit access code be dialed.

primary center—in international terms, a switching center through which trunk traffic is passed and to which local exchanges are connected; i.e., the equivalent of a Class 4 office in the U.S. However, in the U.S. a primary center is defined as a Class 3 office and is used for toll switching. See *class of exchange*.

primary group—group of basic signals that are combined by multiplexing; the lowest level of the multiplexing hierarchy. The term is also used for the signal obtained by multiplexing these basic signals, or for the transmission channel that carries it. Also called *primary block*.

primary rate interface—in ISDN, the interface to the primary rate CCITT 23B + D, 23 channels + 1 signaling channel. See *integrated services digital network*.

primary station—the data terminal in a network that selects and transmits information to a secondary terminal and has the responsibility to insure information transfer.

priority trunk queuing—through a customer-chosen preferential trunk access level, this feature places any caller with this or higher level in the class of service assignment ahead of all callers with a lower trunk access level in the queue of callers waiting for the same trunk group.

privacy and privacy release—all other extensions of a line are unable to enter a conversation in progress on that line unless the initiating station releases the feature.

privacy lockout—automatically splits the connection whenever an attendant would otherwise be included on a call with more than one person. When privacy is provided, the attendant lockout feature is also supplied. A tone warning is generated when the attendant bridges into a conversation in progress.

privacy override—activation of a special pushbutton allows the station user to gain access to a given busy line, even though the automatic exclusion facility is engaged by the station using that line.

private automatic branch exchange (PABX)—See *private branch exchange*.

private automatic exchange (PAX)—dial telephone exchange that provides private telephone service to an organization and that does not allow calls to be transmitted to or from the public telephone network.

private automatic telex exchange (PATX)—exchange used for a private telex network within an organization.

private branch exchange (PBX)—a telephone switch located on a customer's premises that primarily establishes voice-grade circuits, over tie-lines between individual users and the switched telephone network. Typically, the PBX also provides switching within a customer premises' local area and usually offers numerous other enhanced features, such as least-cost routing and call-detail recording.

private digital exchange (PDX)—private exchange employing digital transmission techniques.

private exchange (PX)—exchange serving a particular organization and having no means of connection with a public exchange.

private line—denotes the channel and channel equipment furnished to a customer as a unit for exclusive use, generally with no access to or from the public switched telephone network. Also called *leased line*.

private manual branch exchange (PMBX)—private, manually operated telephone exchange that provides private telephone service to an organization and that allows calls to be transmitted to or from the public telephone network.

Glossary

private network—network established and operated by a private organization or corporation. Compare with *public switched telephone network*.

presentation layer—in the OSI model, that layer of processing that provides services to the application layer, allowing it to interpret the data exchanged, as well as to structure data messages to be transmitted in a specific display and control format.

process—performing operations on information.

process control—the monitoring and controlling of industrial operations by a computer system.

processing, batch—method of computer operation in which a number of similar input items are accumulated and sorted for processing.

processing, line—processing of transactions as they occur, with no preliminary editing or sorting of them before they enter the system.

processor—a computer capable of receiving information, manipulating it, and supplying results.

program—a group of instructions that direct a computer's tasks.

programmable read-only memory—an information storage area that can be recorded by an operator; information stored there can only be altered through special physical processes.

programmer—a person who designs, writes, and tests computer programs.

programming languages—a set of commands used to develop instructions for a computer.

project planning—a preliminary activity that defines areas to be studied and desired objectives before a formal study begins.

PROM—see *programmable read-only memory*.

prompting—messages from a computer that give instructions to the user.

propagation delay—the period between the time when a signal is placed on a circuit and when it is recognized and acknowledged at the other end. Propagation delay is of great importance in satellite channels because of the great distances involved.

protector—interface between inside and outside plant providing against hazardous voltages or currents.

protocol—a set of strict procedures required to establish, maintain, and control communications. Protocols can exist at many levels in one network such as link-by-link, end-to-end, and subscriber-to-switch. Examples include BSC, SDLC, X.25, etc.

protocol conversion—the process of translating the protocol native to an end-user device (e.g., a terminal) into a different protocol (e.g., ASCII to BSC), allowing that end-user device to communicate with another device (e.g., a computer) with which it would otherwise be incompatible. Protocol conversion can be performed by a dedicated device (a protocol converter), by a software package loaded onto an existing system, such as a general-purpose computer, front-end processor, or PBX system, or by a value-added network, such as Telenet.

PRX—program.

PSE—packet switching exchange.

PSK—see *phase-shift keying*.

PSS—British Telecom's Packet Switch Stream network. Packet switched data network operating to X.25 standards.

PSTN—see *public switched telephone network*.

PTAT—private transatlantic fiber optic cables.

PTT—see *Postal, Telegraph, and Telephone Organization*.

pty—party.

public data network (PDN)—digital leased circuits, digital switched circuits, or packet switching network that is designed to provide low error-rate data transmission by using digital rather than analog techniques.

public switched telephone network (PSTN)—the complete public telephone system, including telephones, local and trunk lines, and exchanges.

Public Utilities Commission (PUC)—a state regulatory body.

pulse—a momentary, sharp alteration in the current or voltage produced in a circuit to operate a switch or relay which can be detected by a logic circuit; a sharp rise and fall of finite duration.

pulse amplitude modulation (PAM)—form of modulation in which the amplitude of the pulse carrier is varied in accordance with successive samples of the modulating signal.

pulse carrier—series of identical pulses intended for modulation.

Glossary

pulse code modulation (PCM)—used to convert an analog signal into a digital bitstream for transmission. Digital transmission technique that involves sampling of an analog information signal at regular time intervals and coding the measured amplitude value into a series of binary values, which are transmitted by modulation of a pulsed, or intermittent, carrier. A common method of speech digitizing using 8-bit code words, or samples, and a sampling rate of 8 KHz.

pulse duration modulation (PDM)—see *pulse width modulation*.

pulse width modulation (PWM)—process of encoding information based on variations of the duration of carrier pulses. Also known as *pulse duration modulation (PDM)*.

pulsing—transmission of address information to an exchange by means of digital pulses. Pulsing methods include multifrequency, rotary dial, and revertive.

pushbutton dialing—use of keys or pushbuttons instead of a rotary dial to generate a sequence of digits to establish a circuit connection. The signal form is usually multiple tones. See *dual tone multifrequency signaling (DTMF)*.

pushbutton dialing to stations—special attendant console feature in which the switching system is served by rotary dial exchange trunk circuits. A 10-button keyset is provided on the console to allow fast dialing of extension numbers in order to complete incoming outside calls.

PWM—see *pulse width modulation*.

pwr—power.

PX—see *private exchange*.

PX64—expected to be the videoconferencing standard accepted worldwide by 1990 (for 2-way, full motion video).

Q

Q—queue.

QOS—see *quality of service*.

quad—a cable of four separately insulated conductors, twisted together in such a way as to provide two pairs.

quality assurance—the actions necessary to provide suitable confidence that an item will operate satisfactorily.

quality of service (QOS)—measure of the performance of a telephone system in terms of the quality of the lines and the amount of call blocking experienced.

quantization noise—signal errors caused by the process of digitizing a continuously variable slope.

query—a request for information entered while the computer system is processing.

queue—series of telephone calls, arranged in sequence, the two ends being the head and tail. New calls are added to the tail. Calls can be removed either from the head or tail.

queuing—1) in telephony, a feature that allows calls to be “held” or delayed at the origination switch while waiting for a trunk to become available. 2) sequencing of batch data sessions.

R

RACE—Research and Development Program in Advanced Communications in Europe. European Community sponsored project for the development of broadband technologies.

rack—framework or structure on which apparatus is mounted, usually by means of shelves or mounting plates. Also known as a *bay*.

radio determination—the use of radio waves to determine the position or velocity of, or to collect the information for, a remote object.

radio determination-satellite service (RDSS)—a service using one or more satellites for radio determination.

radio channel—frequency band allocated to a service provider or transmitter.

radio circuit—physical circuit consisting of two unidirectional radio links and connections to terminal exchanges.

radio communications—any telecommunications by means of radio waves.

radio frequency (RF)—a frequency that is higher than the audio frequencies but below the infrared frequencies, usually above 20 KHz. See also *electromagnetic spectrum*.

radio frequency interference (RFI)—signal interference generated at or near a received wavelength. The FCC establishes RFI standards to reduce RFI.

radio paging—provides attendant and station user dial access to customer-owned radio paging equipment to selectively tone-alert or voice-page individuals carrying pocket radio receivers. The paged party can answer by dialing an answering code from a station within the PBX system.

radio paging access with answerback—allows access to customer-provided paging systems and provides the capability in PBX to connect the paged party to the paging party when the former answers the radio page by dialing a special code from any PBX telephone.

Glossary

radio wave—electromagnetic waves of frequencies between 20 KHz and 3 GHz approximately.

RAM—see *random-access memory*.

random access—a storage device that allows information or blocks of information to be read in any order; e.g., a disk.

random-access memory—a storage technique in which the time required to obtain data is independent of the location.

raster—a scanning pattern used in generating, recording, or reproducing television, facsimile, or graphics images on a screen; raster scanning.

rate—charge for a particular service or equipment usage.

rate averaging—telephone companies' method for establishing uniform toll rates based on distance rather than on the relative cost and/or volume of telephone traffic on a given route.

rate base—total invested capital on which a regulated company is entitled to a reasonable rate of return.

rate center—defined geographic point used by telephone companies in determining distance measurements for inter-LATA mileage rates.

rate of return—percentage net profit that a telephone company is authorized to earn.

RBHC—see *Regional Bell Holding Company*.

RCL—see *recall*.

RDT—see *recall dial tone*.

RDY—ready.

read-only memory—a storage location where information is permanently stored and cannot be altered.

read/write head—a device that records and senses magnetic spots on a magnetic disk.

reading—transferring information from a memory, disk, or magnetic tape drive.

real storage—the holding of information so that the central processing unit can obtain and return it directly; synonymous with *processor storage*.

real storage address—the actual physical location of information in memory.

realtime—1) pertaining to actual time during which a physical process transpires. 2) pertaining to an application in which response to input is fast enough to effect subsequent

input, as when conducting the dialog that take place at terminals in interactive systems.

realtime clock—a device in a computer that keeps the time-of-day and makes this information available to programs.

reasonableness checks—tests made on information reaching a realtime system or being transmitted from it to ensure that the data in question lie within a given range.

recall (RCL)—PBX feature allowing a station user engaged in a call to signal the operator, often by a switchhook flash, to enter the conversation.

recall dial tone (RDT)—stutter, or interrupted, dial tone indicating to a station user that the switchhook flash has been properly used to gain access to system features.

receive only (RO)—1) a printer terminal without a keyboard for data entry. 2) a satellite earth station capable of receiving, but not transmitting, a signal.

receiving perforator (reperforator)—telegraph instrument in which the received signals cause the code of the corresponding characters or functions to be punched in a tape.

record—a single, logically associated information group.

record communication—communication that produces a hardcopy record of the transmission, such as teletypewriter and facsimile.

recorded information service—special type of access trunk that, when dialed, will connect the caller to a prerecorded message.

recorded announcement service—special type of access trunk that, when dialed, will connect the caller to a prerecorded message.

recording density—the number of bits that can be written on a specific area of magnetic media; generally measured in bits per inch (bpi).

recovery—the restoration of a system to full operation after a malfunction has been corrected.

recovery from fallback—when the system has switched to a fallback mode of operation and the cause of the fallback has been removed. This is the process that restores the system to its former condition.

red, green, blue monitor (RGB)—a color display screen for computers.

redundancy—1) portion of the total information contained in a message that can be eliminated without loss of essential information. 2) provision of duplicate, backup equipment to immediately take over the function of equipment that

Glossary

fails. 3) in a database, the storage of the same data item or group of items is two or more files.

redundancy check—automatic or programmed check based on the systematic insertion of components or characters used especially for checking purposes.

redundant processor—a computer that duplicates or partially duplicates the operation of another computer to substantially reduce the possibility of system failure.

reed switch—special type of relay consisting of five moving, reed-like contacts controlled by an electromagnet where the reeds themselves are part of the magnetic as well as the electrical circuit being controlled.

reference database—a file of information that directs users to a primary source for additional details or for the complete text; there are two types: bibliographic and referral. Contrast with *source database*; see also *database*, *bibliographic database*, and *referral database*.

referral database—a file of information that contains citations to such nonprint materials as individuals and organizations; see also *database* and *reference database*.

regenerative repeater—1) repeater utilized in telegraph applications to retune and retransmit the received signal impulses and restore them to their original strength. These repeaters are speed- and code-sensitive and are intended for use with standard telegraph speeds and codes. 2) repeater used in PCM or digital circuits which detects, retimes, and reconstructs the bits transmitted. See also *repeater*.

regenerator—see *regenerative repeater*.

Regional Bell Holding Company (RBHC)—one of the seven holding companies formed by the divestiture of AT&T to oversee the 22 Bell Operating Companies (BOCs), which, provide regulated services, and other subsidiaries, which provide nonregulated services; includes Bell Atlantic, NYNEX, BellSouth, Pacific Telesis, U S WEST, Southwestern Bell Corp., and Ameritech.

regional center—see *class of exchange*.

regional (computer) network—a communications pathway with connections limited to a defined geographic area.

register—first unit in the assembly of common control equipment in an automatic exchange. The register receives address information in the form of dial pulses or dual-tone multifrequency (DTMF) signals and stores it for possible conversion or translation.

regression analysis—a technique for determining the mathematical expression that best describes the functional relationship between two or more variables.

reinitiation time—time required for a device or system to restart (usually after a power outage).

relational database—a method of organizing a file of information to link information contained in separate records.

relay—device, operated electrically, that causes by its operation abrupt changes in an electrical circuit, such as breaking the circuit, changing the circuit connection, or varying the circuit characteristics.

release with howler—if a phone stays off-hook without originating a call, the system transmits a loud tone over the line and then “disconnects” it, ignoring it until it goes on-hook again.

relocatability—a capability that allows programs or information to be transferred to different places in main memory at different times without modifying the program.

relocatable addresses—the memory locations used in a group of instructions that can be changed to any location in main storage; see also *main storage*.

remote access—pertaining to communications with a computer or PBX in one location from a device that is physically removed from the location of the computer.

remote-access software—sometimes called remote-control software, this type of program is a superset of the asynchronous communications software market. Remote-access software allows a PC to have complete control over another PC at a different site.

remote batch—transmission of bulk information stored by a distant terminal or computer.

remote computing—using a distant processor to manipulate data.

remote data concentration—communications processors used for multiplexing data from low-speed lines or terminals onto one or more high-speed lines; see also *multiplexing*.

remote file access (RFA)—permits user applications in most languages, as well as most HP-provided utility programs, to access remote files and remote peripherals transparently.

remote file sharing (RFS)—permits multivendor access to data on remote machines.

remote job entry (RJE)—see *remote batch*.

remote maintenance—feature or service in which a service technician can dial the PBX and be connected, usually through the attendant, to the system processor to test or modify the program.

Glossary

remote station lamp field—for use by stations, usually staffed by secretarial personnel, that frequently answer several station lines.

remote traffic measurement—traffic and feature usage data transmitted by the system to a distant service technician.

reorder tone (RT)—tone signal placed on a line by the switching equipment to tell the user that an error has been made in dialing the number or selecting a feature, and/or that the call cannot be completed.

repagination—an automatic routine that changes page endings if text is inserted or deleted within a document or a new page length is desired.

repeater—1) in analog transmission, equipment that receives a pulse train, amplifies it, and retimes it for retransmission. 2) in digital transmission, equipment that receives a pulse train, reconstructs it, retimes it, and then amplifies the signal for retransmission. 3) in fiber optics, a device that decodes a low-power light signal, converts it to electrical energy, and then retransmits it via an LED or laser-generating light source, often including some form of signal amplification. See also *regenerative repeater*.

reperforator—see *receiving perforator*.

reperforator/transmitter (RT)—a teletypewriter unit consisting of a reperforator and a tape transmitter, each independent of the other. It is used as a relaying device and is especially suitable for transforming the incoming speed to a different outgoing speed and for temporary queuing.

replication—the use of a second piece of hardware to reduce the risk of data loss in case of failure; see also *redundancy*, *redundant processor*.

report generator—a group of instructions that creates and prints or displays a selection of stored information.

reprogramming—changing a group of instructions written for one computer so that it will run on another.

request for information (RFI)—general notification of an intended purchase of communications or computer equipment sent to potential suppliers to determine interest and solicit product materials.

request for proposal (RFP)—follow-up to RFI, sent to interested vendors to solicit a configuration proposal, with prices, that meets a user's requirements.

rerun—a repeat of a machine cycle, generally because of a correction, interrupt, or false start.

rerun point—a carefully selected location designed to allow a computer program to repeat a process if it is interrupted.

resale carrier—company that redistributes the services of another common carrier and retails the services to the public.

reserve power—in case of the failure of a PBX power supply, rechargeable batteries can be added to the system allowing system operation for some length of time—15 minutes to 12 hours—after commercial power has failed. Two different situations can occur: 1) the system has enough power to maintain memory but is unable to supply talking voltage and ringing signals, leading to an interruption in services but allowing a fast recovery, and 2) a more adequate reserve power supply will keep the system running during the failure or brownout.

reserved word—a word used in a Cobol source program that has special privileged meaning to the language processor.

reset key—erases all programs in memory and restores a computer to its original mode.

residual error rate, undetected error rate—the ratio of the number of bits, unit elements, characters, or blocks incorrectly received but undetected or uncorrected by the error-control equipment, to the total number of bits, unit elements, characters, or blocks sent.

resolution—measure of the ability of a visual system, e.g., television, facsimile, etc., to reproduce detail. Usually given as resolution along the scanning lines and parallel to the scanning lines, as these two may differ; see also *matrix*.

resource—any means available to system users, including computational power, programs, data files, and storage capacity.

response time—the time period between a terminal operator's completion of an inquiry and the receipt of a response. Response time includes the time taken to transmit the inquiry, process it by the computer, and transmit the response back to the terminal. Response time is frequently used as a measure of the performance of an interactive system.

restriction services—allow the attendant to control the restriction of stations or groups of stations. It can be very useful in hotels, for example, to turn off service to room phones during the time between check out and check in of guests.

return key—when struck, places the position indicator on a display screen at the left margin one line below its previous horizontal position; it is used to end a line of input, a text line, or a paragraph.

reverse-battery signaling—type of loop signaling in which battery and ground are reversed on the tip and ring of the loop to give an off-hook signal when the called party answers.

Glossary

reverse channel—a simultaneous data path in the reverse direction over a half-duplex facility. Normally, it is used for positive/negative acknowledgments of previously received data blocks; see also *half-duplex*.

reverse charge call—call in which the caller specifies that the charge should be paid by the called party. Also called *collect call*.

reverse video—a method of displaying selected symbols on a display screen in a manner that is the opposite of the screen's normal display color.

RF—see *radio frequency*.

RFI—see *radio frequency interference*.

RFI—see *request for information*.

RFP—see *request for proposal*.

RI—1) request for information. 2) radio frequency interference.

ring—1) ring-shaped contact of a plug usually positioned between, but insulated from, the tip and sleeve. 2) audible alerting signal on a telephone line. 3) a network topology in which stations are connected to one another in a closed logical circle, with access to the medium passing sequentially from one station to the next by means of polling from a master station or by passing an access token from one station to another; also called a *loop*.

ring again—the PBX remembers the last number called by a station and will redial it when the feature is activated. Also called *last number dialed*.

ring network—a network topology in which stations are connected to one another in a closed logical circle, with access to the medium passing sequentially from one station to the next by means of polling from a master station, or by passing an access token from one station to another. Also called a *loop*.

ringback tone—see *ringing tone*.

ringdown—to gain the attention of an operator, a ringing current is applied to a line to operate a device producing a steady signal.

ringing key—key whose operation causes the sending of a ringing current.

ringing signal—any AC or DC signal transmitted over a line or trunk for the purpose of alerting a party at the distant end of an incoming call. The signal can operate a visual or sound-producing device.

ringing tone—tone received by the calling telephone indicating that the called telephone is being rung.

ringing transfer—provides for designated bells in a group of stations to ring for incoming calls. Ringing transfer allows additional sets of bells to be designated, with the user controlling which set is to ring.

RJE—remote job entry; see *remote batch*.

RO—see *receive only*.

ROH—receiver off-hook (permanent signal).

rollback—a programmed return to a prior checkpoint or rerun point; see also *checkpoint/restart facility*, *rerun point*.

ROM—see *read-only memory*.

room cutoff—allows guest telephones to be restricted from outgoing calls when the room is unoccupied (hotel check-out). This feature is activated on an individual station basis from the front desk.

room status—provides room status indication from hotel console, with optional printer, for the following conditions:

- Room vacant, • Room occupied, • Room reserved,
- Message registration data, • Message waiting, • Wake-up, and • Do not disturb.

rotary dial calling—system that will accept dialing from conventional rotary dial sets which generate pulses, although pushbutton (DTMF) dial sets offer faster calling and greater reliability.

rotary output to exchange—many systems are equipped to provide pushbutton dial service in all areas. In cases where the telephone exchange trunks are not designed to accept tone signaling, the system will translate the number entered by a station in tones into rotary dial pulses that can be processed by the exchange.

route advance—variation of Automatic Route Selection that allows the caller to select the first-choice trunk group. If that group is busy, the system will attempt to place the call over alternate trunk groups. Unlike ARS, translation is not provided. If ARS is available, Route Advance is generally unnecessary.

routine—a set of general-use instructions designed to accomplish a task.

routing—assignment of the communications path by which a message or telephone call will reach its destination.

Glossary

routing code—address, or group of characters, in the heading of a message defining the final circuit or terminal to which the message has to be delivered. Also called *routing indicator*.

routing indicator—see *routing code*.

routing table—table associated with a network node that states for each message (or packet) destination the preferred outgoing link that the message should use.

RPG—1) report program generator. 2) a programming language supported on most IBM commercial systems.

RPQ—request for price quotation.

RS-232-C—a technical specification published by the Electronic Industries Association that establishes mechanical and electrical interface requirements between computers, terminals, modems, and communications lines.

RS-422-A—an interface standard published by the Electronic Industries Association establishing requirements for single-pathway communications between computers; provides increased capabilities over RS-232C.

RS-423-A—EIA specification for electrical characteristics of unbalanced-voltage digital interface circuits.

RS-449—EIA specification for general-purpose, 37-position and 9-position interface for data terminal equipment (DTE) and data circuit terminating equipment (DCE) employing serial binary data interchange. RS-449-1 includes Addendum 1.

RT—1) reperforator/transmitter. 2) reorder tone.

RTNR—ringing tone no reply.

run—a single and continuous execution of a group of instructions by a computer.

running time—the interval during which a machine is actually processing.

run-unit—the CODASYL (Conference on Data Systems Languages) word for a single application program execution or task.

S

S-100 bus—used as a standard pathway between components in many 8-bit personal computers.

SAA—see *Systems Application Architecture*.

sampling—a statistical procedure where generalizations are drawn from a relatively small number of observations.

satellite communications—the use of geostationary orbiting satellites to relay transmissions from one earth station to another or multiple earth stations.

satellite computer—a programmable machine that relieves a primary processor of such time-consuming operations as compiling, editing, and controlling input/output devices.

SCA—see *short code address*.

scattering—cause of lightwave signal loss in optical fiber transmission. diffusion of a lightbeam caused by microscopic variations in the material density of the transmission medium.

schematic—diagram that details the electrical elements of a circuit or system.

SCPC—see *single channel per carrier*.

scrambler—coding device applied to a digital channel that produces an apparently random bit sequence. A corresponding device is used to decode the channel, i.e., the coding is reversible.

scratch pad memory—a fast, temporary internal storage area used to hold subtotals for various unknowns that are needed for final results.

screening—prevention of electric, magnetic, or electromagnetic fields from escaping or entering an enclosed area by means of a barrier. Also called *shielding*.

scrolling—the vertical movement (up or down) of lines of data displayed on a CRT screen; see also *smooth scrolling*, *jump scrolling*.

SDA—see *source data automation*.

SDLC—see *synchronous data link control*.

SECAM—Sequential Couleur à Mémoire. Color television broadcasting system using 625 picture lines and a 50Hz field frequency, in which the two color-difference signals are transmitted sequentially instead of simultaneously. Developed and used in France, also used in the Soviet Union. See also *NTSC* and *PAL*.

secondary channel—a low-speed channel established on a four-wire circuit over which diagnostic or control information is passed. User data is passed on the primary, high-speed channels of the circuit.

secondary station—on a communications pathway, a terminal device that has been selected to operate under the control of another terminal device.

secretarial intercept—call forwarding of executives' telephones to a secretary/receptionist who can take messages.

Glossary

sectional center—see *class of exchange*.

security—the protection of information against unauthorized access or use.

seek—the movement of an access arm to locate a storage channel on a disk.

seek time—the interval required to position an access arm at a specified position.

segment—an information set that can be placed anywhere in memory and can be addressed relative to a common origin.

selecting—a communication network technique of inviting another terminal device to receive information.

selective calling—ability of the transmitting station to specify which of several stations on the same line is to receive a message.

selective paging to station—an originating station is able to page to specific individual station instruments.

selective ringing—system designed with the capability of ringing only the desired subscriber's telephone on a multi-party line. Ringers tuned to one of five possible frequencies are used to achieve this effect.

self-relocating program—a group of instructions capable of assigning blocks of memory as they are needed.

self-test and fault isolation—most systems include a processor-check capability that allows the controlling computer to test itself and the rest of the system. If a fault is found an alarm light is lit and a message is given on the system printer teletype, if one is provided. This feature also expedites service since the computer can pinpoint faulty equipment, saving diagnostic time.

semiconductor—a substance that can act as a conductor or insulator of electricity depending on its charged state; it acts as an on/off switch signifying binary digits (1s and 0s).

semiconductor memory—a method of storing information using substances that can act as a conductor or an insulator of electricity; see also *semiconductor*.

sender—device that receives address information from a register or routing information from a translator and then outputs the proper routing digits to a trunk or to local equipment. Sender and register functions are often combined in a single unit.

sensor-based system—a configuration made up of devices that measure an external phenomenon.

sequencing—ordering symbols in a series or according to rank and time.

sequential—listed or recorded in numeric order, generally in ascending order.

sequential access—organizing information in a prescribed ascending or descending order to optimize data storage and access; synonymous with serial access.

Sequential Couleur à Mémoire (SECAM)—color television broadcasting system using 625 picture lines and a 50Hz field frequency, in which the two color-difference signals are transmitted sequentially instead of simultaneously. Developed and used in France, also used in the Soviet Union. See also *National Television Systems Committee (NTSC)* and *phase alternate line (PAL)*.

sequential data set—an information group organized on the basis of the physical position of records.

sequential storage—secondary storage where information is arranged in ascending or descending order, generally by item number.

serial—one bit following another over a single pathway.

serial access—see *sequential access*.

serial interface—an interconnection that transmits information bit by bit rather than a whole character at a time; they are much slower and cheaper than parallel interfaces; see also *parallel interface*.

serial transmission—the conveying of a character of information 1 bit at a time; see also *bit*.

series call—operator arranges for a call to return to the console after the extension it was connected to hangs up. Lets the attendant easily connect a caller with a series of inside extensions without the risk of losing the call. Also called *serial call*.

server—a processor that provides a specific service to the network, e.g., a routing server connects nodes and network of like architectures, a gateway server connects nodes and networks of different architectures, etc.

service bureau—a computer organization that offers time-sharing and software services; typical applications include payroll, billing, and bookkeeping; see also *timesharing, on-line services company*.

service code—one or more digits dialed by a customer to access services such as directory enquiries or operator assistance.

service order—request to a telecommunications vendor or carrier for service or equipment.

Glossary

service organization—any company that contracts to provide for and operate computers not owned or leased by the company.

service terminal—equipment needed to terminate the channel and connect to the station apparatus or customer terminal.

serviceability—the ease with which hardware or software failures can be detected, diagnosed, and repaired.

servicing area—1) region surrounding a broadcasting station where signal strength is at or above a stated minimum. 2) geographic area handled by a telephone exchange, generally equivalent to a LATA.

session— 1) a period of time in which an end user engages in dialog with an interactive computer system. 2) layer 5 of the International Organization for Standardization's (ISO) Open Systems Interconnection (OSI) reference model for network architectures.

SF—single frequency.

shared tenant service—see *tenant service*.

shelf—a device used to mount equipment (such as printed circuit boards, power supplies, etc.) in an equipment rack or bay (see rack). A rack can contain several shelves.

SHF—see *super high frequency*.

shielded pair—two insulated wires in a cable wrapped with metallic braid or foil to prevent interference and provide noise-free transmission.

shielding—see *screening*.

short code address (SCA)—those few digits allocated to any frequently dialed number, which when dialed are translated by the exchange into the required full number. See also *abbreviated dialing*.

SHT—short holding time. See *holding time*.

shutdown—the action of making a system unavailable by disabling all terminals, monitoring the completion of transactions in progress, closing all files in an orderly fashion, and terminating all jobs.

sideband—the frequency band on either the upper or lower side of the carrier frequency band within which the frequencies produced by the process of modulation fall. Various modulation techniques make use of one or both of the sidebands, some of which also suppress the carrier frequency.

sidetone—reproduction in a telephone receiver of sounds picked up by the associated microphone. The microphone can pick up either the voice of the speaker or the room noise.

sign-off—see *log-off*.

sign-on—see *log-on*.

signal—a physical, time-dependent energy value used for the purpose of conveying information through a transmission line. Contrast with *noise*.

signal-to-noise ratio (SNR)—the relative power of a signal compared to the power of noise on a line, expressed in decibels (dB). As the ratio decreases, it becomes more difficult to distinguish between information and interference.

signaling—process by which a caller or equipment on the transmitting end of a line informs a particular party or equipment at the receiving end that a message is to be communicated.

significant digit—the integer of a number that represents the largest percentage.

simple mail transfer protocol (SMTP)—service specifically for electronic mail that functions as a unified "post office" for addressing mail to all users on all nodes of both wide area and local networks.

simplex—pertaining to the capability to transmit in one direction only. Contrast with *half-duplex* and *full-duplex*.

simplex circuit—circuit permitting the transmission of signals in one specified direction only.

simscrip—see *simulation languages*.

simulations—the use of programming techniques to duplicate the operation of one computing system on another computing system.

simulation languages—sets of commands, such as Simscript II or GPSS, used for duplicating tasks as they would run on other systems.

simultaneity—the facility of a computer to allow input/output at the same time between the computer and storage devices.

single channel per carrier (SCPC)—transmission system in which a physical channel is allocated solely to one carrier for the duration of the transmission.

single-digit dialing—provides for single-digit dialing to reach a preselected group of stations. A variation of speed dialing, it also helps reduce the need for key systems by replacing the intercom function.

Glossary

single-mode fiber—a fiber with a small core diameter allowing the propagation of a single light path.

single-precision—the number of memory locations used for a number in the computer; each number has one location.

single-sideband transmission—to make efficient use of the frequency band available, the carrier and the unwanted sideband of an amplitude-modulated wave can be filtered out so that only the sideband that contains all the information is transmitted.

single threading—a group of instructions that completes the processing of one message before starting another; see also *multithreading*.

sink—the terminal connection that collects overflow transmissions on a communications pathway.

16-bit system—refers to the number of 0s and 1s in a word that can be processed, stored, and recalled at one time in one machine cycle; longer word lengths increase efficiency and accuracy but also increase complexity.

68000—a 32-bit internal and 16-bit external word length microprocessor developed by Motorola.

6502—an 8-bit microprocessor developed by MOS Technology.

slave—a terminal or computer controlled by another computer.

sleeve—1) third contacting part on a telephone plug preceded in the location by the tip and ring. 2) the sleeve wire is the third control wire of each telephone in an automatic switching office.

slice architecture—a method of using several dedicated subprocessors to provide full processing capabilities.

smalltalk—a set of commands used for the Xerox Star 8010 workstation; the operating environment can be tailored to individual user requirements.

smart terminal—a display terminal that can operate in either conversational or block mode and can support a full range of local editing capabilities.

SMDR—see *station message detail recording*.

smooth scrolling—the continuous vertical movement (up or down) of lines of data displayed on a CRT screen, much in the same manner as a credit roll at the end of a movie.

SNA—see *Systems Network Architecture*.

SNR—see *signal-to-noise ratio*.

softcard—a circuit board manufactured by Microsoft, Inc. that enables an Apple computer to use the CP/M-80 operating system.

soft sectored—a technique used to identify a section of a storage area by numbers written in the storage area itself rather than by fixed location markers.

software—computer instructions that perform common functions for all users as well as specific applications for particular user needs.

software house—a vendor that develops and markets custom computer programs for specific applications.

software library—a collection of programs used to develop user-application programs, including operating and executive systems, high-level languages, and assemblers.

software package—a program or program series sold together; generally sold in machine language form and with user documentation that describes the operation; see also *machine language*.

SOGT—Senior Officials Group-Telecommunications. An ad hoc committee composed of senior Member State and Commission Officials in the European Community to define, oversee, and agree on the technical work of the GAP committees in telecommunications.

solidstate device—electronic pathways made of solid materials, e.g., chips and bubble memories.

SOP—standard operating procedure.

sort—to rearrange information in numeric, alphabetic, or alphanumeric groups.

sorter/reader—see *MICR reader/sorter*.

source—the entry point where information is input.

source data automation—the capturing of information for machine entry.

source database—a group of files containing a full representation of original information; it provides users with quantitative answers without referencing other sources; see also *database*, *numeric database*, *textual-numeric database*, and *full text database*.

source document—the form on which an original transaction was captured; also, an original record with information that is to be converted into machine-readable form.

source language—a set of instructions written by a programmer in the original programming language.

Glossary

source listing—a record of the computer instructions in program language.

source program—a group of instructions that can be read by the computer.

space—the signal (communications channel state) corresponding to a binary zero. The space condition exists when no current flows (current-loop channel) or when the voltage is more positive than +3 volts (EIA RS-232-C channel). In a neutral circuit, it causes the loop to open or causes the absence of a signal; while in a polar circuit, it causes the loop current to flow in a direction opposite to that for mark impulse.

space-diversity reception—to reduce the effects of fading and attenuation, a radio signal is received at more than one site, the sites being separated by a few wavelengths. The signals are then combined, selected, or both.

space-division switching—method for switching circuits in which each connection through the switch takes a physically separate path.

space segment—in a satellite circuit, the satellite itself, and necessary tracking, monitoring, and control functions. Excludes any ground equipment.

SPC—see *stored program control*.

spcl—special.

speakerphone—telephone device that has a speaker-microphone unit allowing hands-free conversation. Also called *speakerset*.

special character—a symbol that is neither numeric nor alphabetic; in Cobol, some examples include , + - * / .) (.

special-purpose computer—a programmable machine designed to handle a restricted class of applications.

specialized carrier—a company that provides value-added or limited communications facilities. Contrast with *common carrier*.

spectrum—continuous range of frequencies, usually wide in extent, within which waves have some specific common characteristics.

speech circuit—circuit designed for the transmission of speech, either analog or encoded, but which can also be used for data transmission or telegraphy.

speed dialing—feature that enables a PBX or PBX station to store certain telephone numbers and dial them automatically when a short (1-, 2- or 3-digit) code is entered. Also called *speed calling*.

split access to outgoing trunks—provides two separate trunk groups for direct outward dialing. The groups can be accessed by dialing the same trunk access code.

split screen—to section the display portion of a video screen into two or more areas, sometimes called windows, so that the different areas can be viewed and compared simultaneously by the user.

splitting—permits the operator to consult privately with one party on a call without the other party hearing.

spooling—temporary storage of batch input and output on disk or tape files until the processor is ready; see also *batch*.

spot beam—a satellite signal that is concentrated on a small geographic area.

SSB—single sideband; see *single-sideband transmission*.

SSN—see *switched service network*.

ST—start (signal to indicate end of outpulsing).

stack—an area in memory for the temporary storage of information; information stored here is not retrieved by address, but rather in chronological order, last in-first out (LIFO).

standalone program—a group of instructions that can be executed independently of the master control program that runs the computer.

standby central control—second control computer that can be provided to direct PBX operations if the primary one fails.

star—a network topology in which each station is connected only to a central station by a point-to-point link and communicates with all other stations through the central station.

StarLAN—a local network design and specification within IEEE 802.3 standards subcommittee, characterized by 1M bps baseband data transmission over two-pair, twisted-pair wiring.

start bit—in asynchronous transmission, a signal used to signify the beginning of the transmission of a character. See also *asynchronous transmission*.

start-stop (signaling)—signaling in which each group of code elements corresponding to an alphabetical signal is preceded by a start signal which serves to prepare the receiving mechanism for the reception and registration of a character and is followed by a stop signal which serves to bring the receiving mechanism to rest in preparation for the reception of the next character. Also known as *asynchronous* or *start-stop transmission*.

Glossary

start-stop transmission—see *asynchronous transmission*.

statement—in programming, an expression or generalized instruction in source language; see also *source language*.

static RAM—random-access memory that requires continuous power but not continuous regeneration to retain its contents.

station—one of the input or output points of a communications system—e.g., the telephone set in the telephone system or the point where the business machine interfaces the channel on a leased private line.

station busy lamps—individual lamps located on the station instrument, providing a visual indication of each station in the system that is busy on an internal call.

station busy override—indicates that preselected stations have the facility to “preempt” busy circuits and override a private conversation.

station call transfer—station user can transfer incoming exchange calls to another station within the system without the need for attendant assistance.

station camp-on—stations can camp-on to a busy extension; the called station is usually notified of the camp-on by a special signal.

station direct station selection (DSS)—station user is able to place a call to any one of a given number of preselected station lines within the PBX system by depressing a single pushbutton on the station or on the auxiliary pushbutton miniconsole that is associated with each station number included in the arrangement.

station equipment—hardware located at a network station. Telephone examples include rotary-dial and pushbutton telephones, key telephones, speakerphones, and IVDTs.

station forced busy—facility that allows a station user to “busy out” the station for temporary periods of time by dialing a special code.

station hunting—routes a call to an idle station line in a prearranged group when the called station line is busy.

station message detail recording (SMDR)—processor-generated records of all calls originated and/or received by a PBX system.

station message registers—message unit information is centrally recorded on a per-station-line basis for each completed outgoing local call made by the station user. Most systems provide for surcharges on station usage and automatically reset the counter after readout.

station message waiting—a “message waiting” light on a station activated by either a button on another station or by a dialed code to the PBX. This feature can alert hotel guests to messages waiting at the front desk, or office workers to messages taken while they were out.

station monitoring—allows selected stations to monitor any other stations within the system.

station override security—on an individual station line basis, designated stations can be “shielded” against the executive busy override facility that is being used by another station.

station rearrangement and change—allows the customer to move stations, change the features and/or restrictions assigned to a station, administer features associated with electronic telephones, and perform search routines on individual stations in order to identify the services provided for that station. Station rearrangements and changes are made on a per-line basis.

station tone ringing—electronic tone ringer or small loudspeaker that transmits an oscillator-created tone.

station transfer security—if a trunk call is transferred from one station to another and the second station does not answer within a predetermined time interval, the trunk call is automatically rerouted to the attendant console.

station-to-station dialing—system feature that allows calling between stations by direct dialing without the need for operator assistance in call completion.

statistical multiplexing—a time-division multiplexing technique in which time slots are dynamically allocated on the basis of need, i.e., slots are allocated to equipment with data to be transmitted.

status information—data about the logical state of a piece of equipment, e.g., a peripheral device reporting its current state to the computer.

STD—subscriber trunk dialing.

step—one operation in a computer program.

step call—allows the attendant or station user, upon finding that the called station is busy, to call a nearby idle station by dialing a single additional digit when the nearby station number has only a different last digit.

step-by-step switch (SXS)—switch that moves in synchronism with a pulse device such as a rotary telephone dial. Each digit dialed causes the movement of successive selector switches to carry the connection forward until the desired line is reached.

Glossary

stop bit—in asynchronous transmission, the quiescent state following the transmission of a character; usually required to be at least 1, 1.42, 1.5, or 2 bit times long.

stop element—last bit of a character in asynchronous serial transmission, used to ensure recognition of the next start element.

storage—a medium used to retain information; it may be internal or external to the computer.

storage capacity—the amount of data that can be contained in an information holding device or main memory, generally expressed in terms of "K" bytes, characters, or words; 1K = 1,024.

storage fragmentation—the inability to assign real storage locations to virtual addresses because the available spaces are smaller than the page size.

storage protection—methods of ensuring the availability and accuracy of mediums where information is retained.

store-and-forward—the process of accepting a message or packet on a communications pathway, retaining it in memory, and retransmitting it to the next station. Synonymous with *message switching*.

stored program computer—a programmable machine that processes information and is controlled by internally retained instructions.

stored program control (SPC)—in processor-controlled switching systems, instructions are held in the form of a program in an electrically alterable store, allowing additions and changes to functions to be made simply by altering the programs.

straightforward outward completion—operator can place an outgoing call for the station user, either by dialing "0" or by an intercept arrangement, without requiring the station user to hang up and redial the operator.

strap—a hard-wired connection. A strapping option is one that is implemented by changing wires.

stress—structural engineering systems solver language; a set of commands used for civil engineering applications including design and sophisticated modeling and analysis; this language has been replaced by STRUDL.

string—a continuous set of characters treated as a single unit.

structured programming—program design and documentation techniques that approach a task as a series of subtasks.

structured walkthrough—a formalized technique allowing a programmer to describe the step-by-step functioning of the program to other programmers.

stunt box—1) device to control the nonprinting functions of a teletypewriter terminal, such as a carriage return and line feed. 2) device to recognize line control characters.

STX—start of text (of message).

subprogram—see *subroutine*.

subroutine—program segments that perform a specific function.

subscriber line—telephone line connecting the exchange to the subscriber's station. Also called *access line* and *subscriber loop*.

subscriber loop—see *subscriber line*.

subscript—a symbol that appears below and after a character; also, programming notation for locating information in a table.

substrate—the physical material on which an electronic pathway in a microprocessor is fabricated.

subsystem—a secondary or subordinate configuration, generally capable of operating within a controlling configuration.

subvoice grade channel—channel with bandwidth narrower than that of voice grade channels. Such channels are usually subchannels of a voice grade line.

super high frequency (SHF)—denotes frequencies from 3 GHz to 30 GHz

supergroup—assembly of five 12-channel groups occupying adjacent bands in the spectrum for the purpose of simultaneous modulation or demodulation; i.e., 60 voice channels.

superscript—a symbol that appears above the character base line.

supervised station release—station that is "off-hook" (that is, the user has not dialed, or is connected to a busy signal for more than a predetermined time interval) is automatically routed to the attendant console.

supervision—process of detecting a change of state between idle and busy conditions on a circuit.

supervisor—the part of an operating system that coordinates the use of resources and maintains the flow of CPU operations.

Glossary

supervisory control—characters or signals that automatically actuate equipment or indicators at a remote terminal.

supervisory lamp—lamp illuminated during a call and indicating to an operator the status of the call.

supervisory program—computer program designed to coordinate, service, and augment the machine components of the system and coordinate and service application programs. They handle work scheduling, input/output operations, error actions, and other functions. Also known as *operating system*. See also *application program*.

supervisory signal—1) signal that indicates whether a circuit is in use or that gives an indication of status or change of status in a telephone system. 2) signal used to indicate the various operating states of circuit combinations.

supplier—see *vendor*.

swap—in systems using timesharing, to write the main storage elements of a job to auxiliary storage and read another job into main storage; see also *timesharing*, *main storage*, and *auxiliary storage*.

swapping—see *swap*.

switchboard—equipment on which switching operations are performed by operators or attendants.

switched line—one of a series of lines that can be interconnected through a switching center; a line on the public telephone network. Contrast with *leased line*.

switched loop operation—attendant position is arranged so that each call requiring attendant assistance is automatically switched to one of several switched loops in position. Normally, the call automatically releases from the position when answered by the called station (released loop operation). Incoming calls are queued in the order of arrival when all attendant positions are busy and are switched to each attendant position automatically to distribute the call load evenly to each attendant. A console lamp indication is normally given to the attendant when calls are waiting in the queue to be served.

switched message network—a network service, such as Telex, providing interconnection of message devices such as teletypewriters.

switched network—a multipoint communications pathway with circuit-switching capabilities; e.g., the telephone network, Telex, and TWX.

switched service network (SSN)—network consisting of terminals, transmission links, and at least one exchange, on which any user can communicate with any other user at any time.

switched virtual call (SVC)—the stream of packets which forms the data flow for a single data message.

switchhook—switch on a telephone set, associated with the structure supporting the receiver or handset, and often used to signal the switching equipment or an attendant during a call, e.g. to transfer the call.

switching—establishment of transmission path from a particular inlet to a particular outlet of a group of such inlets and outlets.

switching center—location that terminates multiple circuits and is capable of interconnecting circuits or transferring traffic between circuits.

switching matrix—see *matrix*.

switchover—when a failure occurs in the equipment, a switch can occur to an alternative component.

SXS—step-by-step switch.

symbolic language—a set of commands that uses mnemonic terms that are easy to remember; see also *mnemonic*.

symbolic name—a data field identifier that the computer changes into storage addresses; Fortran and Basic call this a variable.

sync character—a symbol or defined bit pattern that is used by the receiving terminal to adjust its clock and achieve synchronization.

synchronization—the process of adjusting a receiving terminal's clock to match the clock of the transmitting terminal.

synchronous—having a constant time interval between successive bits, characters, or events. Synchronous transmission uses no redundant information (such as the start and stop bits in asynchronous transmission) to identify the beginning and end of characters, and thus is faster and more efficient than asynchronous transmission. The timing is achieved by transmitting sync characters prior to data; usually, synchronization can be achieved in two or three character times.

synchronous communications—high-speed transmission of contiguous groups of characters; the stream of monitored and read bits is using a clock rate.

synchronous data link control (SDLC)—an IBM communications line discipline or protocol associated with SNA. In contrast to BSC, SDLC provides for full-duplex transmission and is more efficient.

synchronous network—network in which all the communications links are synchronized to a common clock.

Glossary

synchronous transmission—transmission process whereby the information and control characters are sent at regular, clocked intervals so that the sending and receiving terminals are operating continuously in step with each other.

syntax error—a system response to a mistake in instruction, such as a transposition of characters or an omission of a character or word.

sysgen—see *system generation*.

syslog—see *system log*.

system—an organized collection of parts or procedures designed to perform a function; generally refers to the processor, peripherals, and software.

system control programming—vendor-supplied groups of instructions that are fundamental to the operation and maintenance of a configuration.

system design—the specification of the working relations between all the parts of a configuration.

system generation—(SYSGEN); the process of using a master control program to assemble and link all the parts that constitute another operating system.

system library—a collection of data sets in which the various parts of a master control program are stored.

system log—a record of job-related information, operational data, descriptions of unusual occurrences, commands, and messages to or from the operator.

system programmer—a technical expert who plans, generates, maintains, extends, and controls program development and implementation.

system reliability—the probability that a configuration will perform its specified task properly under stated environmental conditions.

system resource—any facility of a computing configuration that can be allocated to a task.

system software—see *systems and support software*.

systems analyst—an individual who defines application problems, determines system specifications, recommends equipment changes, designs data processing procedures, and devises data verification methods; also, one who prepares block diagrams and record layouts from which the programmer prepares flowcharts; may assist with or supervise the preparation of flowcharts.

systems and support software—the variety of computer instructions and associated tools including assemblers, com-

pilers, subroutine libraries, operating systems, and application programs generally provided by the computer vendor.

Systems Application Architecture (SAA)—IBM's announced framework for allowing development of consistent applications across six software environments (TSO/E, CICS/MVS, IMS/ESA/TM, VM/CMS, OS400, and OS/2 Extended Edition) running on three hardware computing platforms (System/370 ESA, OS/400, and PS/2).

systems approach—a general term for reviewing all implications of a condition or group of conditions rather than the narrow implications of a problem at hand.

Systems Network Architecture (SNA)—IBM's standardized relationship between its virtual telecommunication access method (VTAM) and the network control program (NCP/VS).

systems test—a complete simulation of an actual running configuration for purposes of ensuring the adequacy of the configuration.

T

T carrier—a time-division multiplexed, digital transmission facility, operating at an aggregate data rate of 1.544M bps and above. T carrier is a pulse code modulation (PCM) system using 64K bps for a voice channel.

T,R—1) tip, ring. 2) transmit and receive.

T1—a digital carrier facility used to transmit a DS1 formatted digital signal at 1.544M bps; the equivalent of 24 voice channels. The European equivalent transmits at 2.048M bps.

T1C—a digital carrier facility used to transmit a DS1C formatted digital signal at 3.152M bps; the equivalent of 48 voice channels.

T2—a digital carrier facility used to transmit a DS2 formatted digital carrier signal at 6.312M bps; the equivalent of 96 voice channels.

T3—a digital carrier facility used to transmit a DS3 formatted digital carrier signal at 44M bps; the equivalent of 672 voice channels.

T4—a digital carrier facility used to transmit a DS4 formatted digital carrier signal at 273M bps; the equivalent of 4,032 voice channels.

TAAS—trunk answer from any station.

table-driven—describes a logical computer process, widespread in the operation of communications devices and networks, where a user-entered variable is matched against

Glossary

an array of predefined values. Also, a frequently used logical process in network routing, access security, and modem operation.

table look-up—searching for information in an ordered collection of data; also, the process of using an ordered collection of data stored in main memory during the running of a group of instructions to obtain a function value.

TACS—total access communications system. A derivative of the AT&T-developed analog cellular radio standard AMPS (advanced mobile phone service) adopted by U.K. Cellnet and Racal-Vodafone, operating at 900 MHz.

tail—in satellite networks, any terrestrial extension used to connect end users to an earth station.

talking battery—dc voltage supplied by the exchange to the subscriber's loop to operate the carbon transmitter in the handset.

talking path—in a telephone circuit, the transmission path consisting of the tip and ring conductors.

tandem—the connection of networks or circuits in series, i.e., the connection of the output of one circuit to the input of another.

tandem data circuit—a data circuit that contains two or more data circuit terminating equipment (DCE) in series.

tandem exchange—exchange that serves to switch traffic between other exchanges when direct trunks are not available. Also called *tandem office*.

tandem office—see *tandem exchange*.

tandem switching—switching of circuits between exchanges only.

tandem tie line switching—PBX permits tie lines to "tandem" through the switch. This means that an incoming tie line call from the distant PBX receives a dial tone instead of automatically connecting with the operator. The outgoing line can be a local trunk or another tie line that links a third system.

tandem trunk—circuit between a tandem exchange and an exchange.

tape drive—a mechanism for controlling the movement of magnetic storage material past a reading or writing head.

tariff—the published rate for the use of a specific unit of equipment, facility, or type of service provided by a communications common carrier; also, the vehicle by which the regulating agencies approve or disapprove such facilities or services; see also *common carrier*.

TAS—see *telephone answering service*.

TASI—see *time-assignment speech interpolation*.

task—a unit of work.

TAT—transatlantic telephone cable.

TCAM—see *telecommunications access method*.

TCM—see *traveling class mark*.

TCP/IP—see *transmission control protocol/internet protocol*.

TDF—trunk distribution frame. See *main distribution frame*.

tdm—tandem.

TDM—see *time-division multiplexing*.

TDMA—see *time-division multiple access*.

Telco—telephone company.

telecommunication lines—telephone and other communication pathways that are used to transmit information from one location to another.

telecommunications—any process that permits the passage of information from a sender to one or more receivers in any usable form (printed copy, fixed or moving pictures, visible or audible signals, etc.) by means of any electromagnetic system (electrical transmission by wire, radio, optical transmission, waveguides, etc.). Includes telegraphy, telephony, video-telephony, data transmission, etc.

telecommunications access method (TCAM)—an IBM macro language for creating communication applications programs and message control.

telegram—hard copy information, whether in written, printed, or pictorial form, sent to the general telegraph service for transmission and delivery to the addressee.

telegraph channel—transmission media and intervening apparatus involved in the transmission of telegraph signals between two terminal sets or two intermediate telegraph installations.

telegraphy—branch of telecommunications concerned with processes providing reproduction, at a distance, of written, printed, or pictorial matter or the reproduction at a distance of any kind of information in such form.

teleinformatics/telematics—terms describing nonvoice services, particularly those emerging from the integration of computing and telecommunications.

Glossary

telephone answering service (TAS)—private concern that answers telephone calls and takes messages for a large number of different people and organizations. Because the called number is identified on a special console, the answering service attendant can answer each call as if he or she were actually on the called party's premises.

telephone channel—transmission path designed for the transmission of signals representing human speech or other telephone communication (e.g., facsimile) requiring the same bandwidth. The bandwidth allotted is usually 64K bps, but can be reduced to 32K or even 16K bps with multiplexing techniques.

telephone circuit—electrical connection permitting the establishment of a telephone communication in both directions between two points. Also called *telephone line*.

telephone exchange—switching center for interconnecting the lines that terminate therein. Also called *central office (CO)*.

telephone frequency—any frequency within that part of the audiofrequency range essential for the transmission of speech of commercial quality, i.e., 300 to 3000Hz. Also called *voice frequency (VF)*.

telephone line—see *telephone circuit*.

telephone receiver—device within the handset that converts electrical energy into sound energy and designed to be placed next to the ear.

telephony—generic term describing voice telecommunications.

teleprinter—see *teletypewriter (TTY)*.

teleprocessing—the processing of information that is received from or sent to remote locations by way of telecommunication lines; such systems are necessary to hook up remote terminals or connect geographically separated computers; see also *telecommunications*.

teletext—generically, one-way data transmission designed for widespread broadcasting of graphics and textual information for display on subscriber televisions or low-cost terminals.

teletype—frequently used as a generic name for keyboard/printers and for asynchronous transmission.

teletypewriter (TTY)—start-stop apparatus comprising a keyboard transmitter, together with a printing receiver.

television receive-only—(TVRO) an antenna designed specifically to receive but not transmit a television signal.

telex—1) dial-up telegraph service enabling subscribers to communicate directly and temporarily among themselves by means of start-stop apparatus and of circuits of the public telegraph network. The service operates worldwide using Baudot equipment. 2) Western Union provides such services within the U.S. and abroad under its Telex and Telex II trademarks. Other international record carriers (IRCs) also provide telex services to both the domestic and international markets.

Telex II—see *TWX*.

TELNET (TELEtype NETWORK)—provides a virtual terminal capability for accessing remote systems as a terminal.

temporary station disconnection—allows the attendant to completely remove selected stations from total service at any time on a temporary basis.

temporary storage—main memory locations reserved for intermediate programming results.

tenant service—two or more closely located customers can simultaneously be served by the same PBX equipment. Each customer is provided with separate attendant facilities, dedicated trunk facilities, and separate feature and class of service complements. It is often provided by the owner of an office building to the tenants.

terminal—1) a point at which information can enter or leave a communication network. 2) any device capable of sending and/or receiving information over a communication channel.

terminal emulation—imitation of a specific terminal (VT100, for example) by a device, such as a PC, through software. PCs often use terminal emulation methods to connect to specific hosts, such as Digital VAXs or IBM mainframes, with which they would otherwise be unable to communicate.

terminal job—in systems with timesharing, the processing done on behalf of one terminal user from log-on to log-off; see also *timesharing*.

terminal session—see *session*.

terminal user—in systems with timesharing, anyone who is eligible to log-on; see also *timesharing*.

test center—facility that receives customer trouble reports, tests communications lines and equipment, and can dispatch repair technicians.

test data generator—communications instructions for forming files containing sets of information developed specifically to ensure the adequacy of a computer run or system.

Glossary

test desk—switchboard equipped with testing apparatus, so arranged that connections can be made from it to telephone lines or exchange equipment for testing purposes. Also called *test board*.

test generators—a software aid used to help ensure the adequacy of new programs.

testing (program)—the process of executing a group of instructions with the intention or goal of finding errors.

testing (system)—the verification and/or validation of the configuration; it is a verification process when done in a simulated environment and a validation process when performed in the actual environment.

text—the information portion of a transmitted message, as contrasted with the header, check characters, and end-of-text characters.

text editing—manipulation of the information portion of a transmitted message, by changing or rearranging symbols, words, sentences, and paragraphs.

textual-numeric database—a collection of files containing records made up of a combination of data elements; see *database, source database, numeric database, full-text database*.

TFR—transfer.

TGB—trunk group busy.

TGW—trunk group warning.

The Source—a commercial information utility that provides a wide variety of information that can be accessed remotely.

third-party maintenance—see *service organizations*.

3780—a batch protocol used to communicate with an IBM mainframe or compatible system; see also *batch, protocol*.

32-bit system—refers to the number of 0s and 1s making up a word that can be processed, stored, and recalled at one time in a single machine cycle.

3270—IBM's interactive communications terminal standard used to communicate with an IBM mainframe or compatible systems.

thrashing—in simulated storage systems, a condition in which a system can do little useful work because of excessive transfer of programs in and out of memory.

three-way conference transfer—by depressing the switchhook a user can dial another extension and either 1) hang up and transfer call, 2) get information from the called

party and then resume the first call, or 3) bridge all three parties together for a three-way conference.

throughput—the total useful information processed or communicated during a specified time period; expressed in bits per second or packets per second.

tie line (TL)—private-line communications channel of the type provided by communications common carriers for linking two or more points together, typically PBXs. Also called tie trunk; see also *common carrier, trunk line, leased line*.

tie trunk access—allows the system to handle tie lines that can be accessed either by dialing a trunk group access code or through the attendant. Tie lines link a PBX or a distant key system.

time-assignment speech interpolation (TASI)—specialized switching equipment that connects a party to an idle circuit while speech is taking place, but disconnects the party when speech stops, so that a different party can use the same circuit. During periods of heavy traffic, TASI can improve line efficiency from 45 percent to 80 percent.

time-division multiple access (TDMA)—communicating devices at different geographical locations share a multipoint or broadcast channel by means of a technique that allocates different time slots to different users.

time-division multiplexing (TDM)—means of obtaining a number of channels over a single path by dividing the path into a number of time slots and assigning each channel its own intermittently repeated time slot. At the receiving end, each time-separated channel is reassembled. Contrast with *frequency-division multiplexing*.

time-division switching—switching method for a TDM channel requiring the shifting of data from one slot to another in the TDM frame. The slot in question can carry a bit, byte, or, in principle, any other unit of data.

time out—a set time period for waiting before a terminal system performs some action. Typical uses include a poll release (when a terminal is disconnected if the time-out period elapses before keying resumes), or an access time-out (when a terminal on a local area network using a CSMA/CD access method is prevented from transmitting for a specified time period after a collision occurs).

timed recall—the PBX can be instructed to place a call at a designated time. When the time comes, the PBX rings the station. When the station answers, the call is placed automatically.

timed recall on outgoing lines—outgoing trunk calls can be automatically transferred to the attendant after a selected time interval. A warning tone is sent to the calling party several seconds before the transfer takes place.

Glossary

time series—a sequence of quantitative data assigned to specific intervals.

timesharing—method of operation in which a computer facility is shared by several users for different purposes at (apparently) the same time. Although the computer actually services each user in sequence, the high speed of the computer makes it appear that the users are all handled simultaneously.

timesharing option (TSO)—an IBM programming system for implementing an environment that is shared by several users simultaneously; it runs under the operating system.

time slice—an interval on the central processing unit allocated for performing a task; once the interval is expired, CPU time is allocated to another task; see also *timesharing*.

tip—contacting part at the end of a telephone plug or the top spring of a jack. The conductors associated with these contacts. The other contact is called a ring.

TL—see *tie line*.

TLF—trunk link frame (crossbar switching).

TLP—see *transmission level point*.

token bus—a local network access mechanism and topology in which all stations actively attached to the bus listen for a broadcast token or supervisory frame. Stations wishing to transmit must receive the token before doing so; however, the next physical station to transmit is not necessarily the next physical station on the bus. Bus access is controlled by preassigned priority algorithms.

token passing—a local area network access technique in which participating stations circulate a special bit pattern that grants access to the communications pathway to any station that holds the sequence; often used in networks with a ring topology.

token ring—a local network access mechanism and topology in which a supervisory frame or token is passed from station to station in sequential order. Stations wishing to gain access to the network must wait for the token to arrive before transmitting data. In a token ring, the next logical station receiving the token is also the next physical station on the ring.

toll call—call outside the local exchange area, charged at toll rates.

toll center—class 4/primary telephone exchange where time- and distance-based toll charge information is collected.

toll-connecting trunk—trunk used to connect a class 5/local exchange to the long-distance network.

toll restriction—blocking a telephone user's access to the toll network.

toll terminal access—allows guest stations to access toll calling trunks. These can be direct dial or operator access depending on the servicing public exchange office.

tone ringing—either a steady or oscillating electronic tone is provided at the station instrument to provide incoming calls with audible signaling. Also called *tone calling*.

tone signaling—transmission of supervisory, address, and alerting signals over a telephone circuit by means of tones.

tone-to-dial-pulse conversion—converts DTMF signals to dial pulse signals when the trunks associated with outgoing trunk calls are not equipped to receive tone signals. Auxiliary dial pulse conversion equipment is not necessary.

TOP—technical office protocols developed by Boeing.

topology—the logical or physical arrangement of stations on a network in relation to one another. See *bus*, *ring*, *star*, and *tree*.

touch sensitive—refers to the technology that enables a system to identify a point of contact on the screen by coordinates and transmit that information to a program.

Touch-tone—registered AT&T trademark for pushbutton dialing. See also *dual tone multifrequency signaling (DTMF)*.

trace packet—in packet switching, a special kind of packet that functions as a normal packet but causes a report of each stage of its progress to be sent to the network control center.

track—a storage channel on a disk or tape made up of a series of sectors.

tracking, telemetry, control, and monitoring (TTC&M)—specialized ground stations used to track and control satellites and to monitor their performance.

traffic—1) messages sent and received over a communications channel. 2) quantitative measurement of the total messages and their length, expressed in hundred call seconds (CCS) or other units.

traffic data to customer—customer can poll switching locations on a daily or hourly basis to obtain traffic measurements, including usage and overflow data. Summary reports, exception reports, and complete traffic register outputs can be obtained.

traffic flow—measure of the density of traffic, expressed in Erlangs.

Glossary

traffic matrix—matrix of which the X,Y element contains the amount of traffic originated at node X and destined for node Y. The unit of measurement could be calls or packets per second, for example, depending on the kind of network.

traffic monitor—PBX feature that provides basic statistics on the amount of traffic handled by the system.

traffic overflow—occurs when traffic flow exceeds the capacity of a particular trunk group and is automatically switched over to another trunk group.

traffic service position system (TSPS)—stored-program computer with telephone operator consoles permitting calls needing operator intervention to be handled as efficiently as possible.

trailer label—information written after a file has been processed to indicate how many logical records make up the file; synonymous with detail file.

train time—initialization time for full-duplex operation on a modem.

transaction file—a collection of records for any business activity or request that is entered into a computer system.

transaction processing—a procedure in which files are interactively updated and results are generated immediately as a result of data entry.

transceiver—device that can transmit and receive traffic.

transducer—a device for converting signals from one form to another, such as a microphone or a receiver.

transfer, all calls—allows the station to make one transfer of an outside line to another internal station.

transfer rate—the speed at which information can be sent across a bus or communications link.

transistor-transistor logic (TTL)—a type of signaling, in which a nominal +5V is equated with logic 1, and a nominal 0V is equated with logic 0.

translation—when using automatic route selection or trunk-to-trunk connection features, the PBX can add or delete area codes and toll access digits from number codes, and toll access digits from numbers so that the call will be handled properly by the switching network.

translator—device that converts information from one system of representation into equivalent information in another system of representation. In telephone equipment, it is the device that converts dialed digits into call routing information.

transmission—sending information in the form of electrical signals over electric wires, waveguides, or radio.

transmission control protocol/internet protocol (TCP/IP)—ensures packets of data are delivered to their destinations in the sequence in which they were transmitted.

transmission level point (TLP)—any point in a transmission system at which the power level of the signal is measured.

transmission speed—the rate at which information is passed through communications lines, generally measured in bits per second (bps).

transmit—to send information from one location to another.

transparency—a transmission mode in which control character recognition is suspended, allowing any bit pattern to be transmitted. The rules for control code recognition are changed so that commands involving specific operations, for example, Escape sequences, are not acted upon, but simply transmitted as data. Certain commands remain active in order to exit transparency mode.

transponder—receiver-transmitter combination that retransmits the received signal greatly amplified and at a different frequency. Communications satellites usually contain several transponders.

transport layer—in the OSI model, the network processing entity responsible, in conjunction with the underlying network, data link, and physical layers, for the end-to-end control of transmitted data and the optimized use of network resources.

traveling class mark (TCM)—when automatic route selection (ARS) or uniform numbering/automatic alternative routing selects a tie trunk to a distant tandem PBX, the traveling class mark is sent over the tie trunk. It is then used by the distant system to determine the best available facility consistent with the user's calling privileges. The TCM indicates the restriction level to be used based on the station, trunk, or attendant originating the call or the authorization code, if dialed.

traveling wave tube amplifier (TWTA)—a satellite component used to amplify an incoming signal before transmitting to receiving earth stations.

tree—a type of bus network topology in which the medium branches at certain points along its length to connect stations or clusters of stations; also called a *branching bus*.

TRS-80—popular family of microcomputers manufactured by Tandy Corporation and sold through Radio Shack outlets.

Glossary

truncation—the removal of one or more digits, characters, or bits from an item of data when a string length or precision of a target variable has been exceeded; also, to cut off a computational procedure at a specified spot.

trunk—transmission paths that are used to interconnect exchanges in the main telephone network. Also, a telephone exchange line that terminates in a PBX.

trunk answer from any station (TAAS)—night service facility activated by the attendant, whereby incoming calls, normally directed to the attendant, activate a common alerting signal on the customer's premises. These calls can be answered by dialing a special single digit from nonrestricted stations.

trunk code—code consisting of one or more digits used to designate a called numbering plan area.

trunk exchange—exchange to which trunk circuits are connected, but not subscribers' lines.

trunk group—those trunks that connect two points, both of which are exchanges and/or individual message distribution points and both of which employ the same multiplex terminal equipment. Also, a discrete group of trunk lines with a specific function in a PBX.

trunk group busy (TGB) indication on attendant position—light associated with a trunk group, activated when all lines in the group are busy. This allows the attendant to easily monitor the status of the system and its line.

trunk group warning (TGW) indication on attendant position—provides the attendant with a visual indication when a certain number of trunks in a trunk group are busy.

trunk prefix—one or more digits to be dialed before the trunk code when making a call to a subscriber outside the local area but in the same country.

trunk reservation—attendant can hold a single trunk in a group and then extend it to a specific station.

trunk-to-tie trunk connections—ability of the switching system to provide the attendant with the capability of extending an incoming exchange call to a tie trunk that terminates within the system.

trunk-to-trunk connections, attendant—provides attendant capability to make all types of trunk connections.

trunk-to-trunk connections, station—station already in connection with an incoming or outgoing trunk is able to utilize the consultation hold and add-on conference circuitry to effect a conference with another circuit or with another trunk.

trunk-to-trunk consultation—allows a station connected to an outside trunk circuit to gain access to a second outside trunk circuit for "outside" consultation; however, no conference capability is available with this feature.

trunks, direct termination—each incoming or combination trunk appears on a key at a console or at a jack position on a cord switchboard and the attendant is always in full visual supervision of the status of all such circuits.

TSO—see *timesharing option*.

TSPS—see *traffic service position system*.

tst—test.

TTC&M—tracking, telemetry, control, and monitoring.

TTL—see *transistor-transistor logic*.

TTN—tandem tie trunk network. Private network making use of tandem switching.

TTY—see *teletype*.

TTY—see *teletypewriter*.

tuning—the process of adjusting computer system control variables to make a system divide its resources most efficiently for a workload.

turnaround time—the time required to reverse the direction of transmission, e.g., to change from receive mode to transmit mode in order to acknowledge on a half-duplex line. When individual blocks are acknowledged, as is required in certain protocols (e.g., IBM BSC). The turnaround time has a major effect on throughput, particularly if the propagation delay is lengthy, such as on a satellite channel.

turnkey system—complete communications system, including hardware and software, assembled and installed by a vendor and sold as a total package.

turtle graphics—a system of language commands developed by Seymour Papert and the Logo group at MIT.

TVRO—see *television receive-only*.

TVS—trunk verification by station. See *trunk verification by customer*.

2780—a batch standard used to communicate with IBM mainframes or compatible systems; see also *batch*.

twisted pair—two insulated wires twisted together but not covered with an outer sheath.

Glossary

two-party station service—PBX system with two internal stations and with selective ringing to each.

two-way alternate operation—see *half-duplex*.

two-way simultaneous operation—see *full-duplex*.

two-way splitting—attendant is able to consult privately with either party (internal or external) on a call.

two-wire channel—a circuit containing two single wires in pair (or logical equivalent) for nonsimultaneous (i.e., half-duplex) two-way transmission. Contrast with *four-wire channel*.

two-wire circuit—a circuit formed by two conductors insulated from each other that can be used as either a one-way transmission path, a half-duplex path, or a duplex path. Contrast with *four-wire circuit*.

TWTA—see *traveling wave tube amplifier*.

TWX—an obsolete but still popular term for Western Union's switched teleprinter exchange service, now called Telex II. Interconnection is provided with the Telex network through Western Union's computers even though speeds and codes are different.

TXD—telephone exchange, digital. Exchange making use of digital transmission techniques.

TXE—telephone exchange, electronic. Exchange making use of electronic switching techniques, as opposed to electromechanical means such as crossbar or step-by-step switches.

TXC—telephone exchange, crossbar. See *crossbar switch*.

TXK—telephone exchange, crossbar. See *crossbar switch*.

TXS—telephone exchange, stronger. See *step-by-step switch (SXS)*.

U

UCD—see *uniform call distribution*.

UDLC—Universal Data Link Control; a bit-oriented communications standard developed by Sperry. (Sperry is now Unisys.)

UG—underground.

UHF—see *ultra high frequency*.

ultra high frequency (UHF)—portion of the electromagnetic spectrum ranging from about 300 MHz to about 3 GHz. The frequency band includes television and cellular radio frequencies.

unattended operation—transmission and/or reception that is controlled automatically and does not require a human operator.

unbalanced-to-ground—with a two-wire circuit, condition in which the impedance-to-ground on one wire is measurably different from that on the other. Compare with *balanced-to-ground*.

unbundled—services, programs, and training that are sold separately from the computer hardware by the manufacturer.

undetected error rate—see *residual error rate*.

uniform call distribution (UCD)—allows calls coming in on a group of lines to be assigned stations as smoothly as possible so that all stations can handle similar loads. Most call distribution systems also provide for a queuing of incoming calls with the longest holding time presented for service first.

uniform numbering plan—permits station users at a PBX to place calls over tie trunks using a uniform dialing plan.

uniform-spectrum random noise—noise distributed over the spectrum in such a way that the power-per-unit bandwidth is constant. Also called *white noise*.

uninterruptible power supply (UPS)—usually includes an inverter, drawing its power from batteries, which generates an extremely "well-behaved" AC power signal for a PBX or other equipment. The UPS cost is related to the amount of power needed and the length of time it must operate during a failure. If a particularly heavy demand is anticipated, the system can be coupled with an auxiliary generator that is started when commercial power is interrupted.

universal data link control (UDLC)—a bit-oriented protocol based on HDLC developed by Sperry Univac. (Sperry is now Unisys.)

universe—the population from which samples are drawn.

UNIX—a multiuser operating system developed by Bell Laboratories.

unlisted number—telephone number that is not listed in the telephone directory and not provided by directory-assistance operators. There is usually an additional charge to the subscriber for the deletion from the directory.

unrestricted extension—PBX extension permitted to make exchange line calls without the assistance of the PBX operator.

update—to modify a master file with current transaction information according to a specified procedure.

Glossary

uplink—the portion of a satellite circuit extending from an earth station to the satellite. Compare to downlink.

UPS—see *uninterruptible power supply*.

USASCII—see *American Standard Code for Information Interchange (ASCII)*.

USITA—United States Independent Telephone Association.

USOC—Uniform Service Order Code.

V

VA—volt ampere.

VAC—volts alternating current.

vacant code intercept—routes all calls made to an unassigned “level” (first digit dialed) to the attendant, a busy signal, a “reorder” signal, or a recorded announcement.

vacant number intercept—usually routes all calls of unassigned numbers to the attendant or a recorded announcement.

VADS—value added data network.

validation—an attempt to find errors by executing a program in a given environment.

validity check—verification that each element of information is an actual character of a code in use.

value-added carriers—vendors that design and enhance the phone network and resell the service.

value-added common carrier—company that sells services of a value-added network. It can be a PTT or subsidiary or an independent company.

value-added network (VAN)—a public data communications network that provides basic transmission facilities (generally leased by the VAN vendor from a common carrier) plus additional, “enhanced” services such as computerized switching, temporary data storage, protocol conversion, error detection and correction, electronic mail service, etc.

value-added service—a communications facility utilizing communications common carrier networks for transmission and providing enhanced extra data features with separate equipment; such extra features, including store-and-forward message switching, terminal interfacing, and host interfacing, are common.

VAN—see *value-added network*.

VANS—value-added network service.

variable—a quantity that can assume any values; see also *field*.

variable-length record—a group of related data fields with independent lengths.

variance—the difference between the expected (or planned) and the actual.

VDC—volts direct current.

VDT—1) video display terminal. 2) visual display terminal; see *cathode ray tube*.

VDU—visual display unit.

vendor—a company that supplies products and services.

Venn diagram—an illustration that represents sets as circles or ellipses to give a graphic representation of basic logic operations.

verification—an attempt to find errors by executing a program in a test or simulated environment.

vertical redundancy check—an error-checking method that uses a parity bit for each character; see also *parity*.

very high frequency (VHF)—portion of the electromagnetic spectrum with frequencies between about 30 MHz and 300 MHz. Operating band for radio and television channels.

very large scale integration (VLSI)—see *large scale integration (LSI)*.

very low frequency (VLF)—refers to frequencies below 30 KHz.

very small aperture terminal (VSAT)—an earth station with a small antenna, usually 6 meters or less. VSATs are typically used in point-to-multipoint data networks.

VF—see *voice frequency*.

VHF—see *very high frequency*.

video—the information displayed on the screen of a CRT.

videoconferencing—an electronic meeting in which geographically separated groups communicate using interactive audio and video technology.

video disk—a rigid, random access storage medium for analog or digital information written and/or read by a laser; see also *optical disks*.

video signal—signal comprising frequencies normally required to transmit pictorial information (1 MHz to 6 MHz).

Glossary

videotex—an interactive data communications application broadcast over the public switched telephone network and received on adapted television receivers. It is designed to allow unsophisticated users to converse with remote databases, enter data for transactions, and retrieve textual and graphics information. Also called *viewdata*.

viewdata—see *videotex*.

virtual address—a memory location that refers to simulated storage and must, therefore, be translated into a real storage memory location when it is used.

virtual address area—the portion of simulated storage where memory locations are greater than the highest memory location of the real memory location area.

virtual address space—the simulated storage assigned to a job, terminal user, or system task.

virtual circuit—proposed CCITT definition for a data transmission service in which the user presents a data message for delivery with a header of a specified format. The system delivers the message as though a circuit existed to the specified destination. One of many different routes and techniques could be used to deliver the message, but the user doesn't know which is employed. Contrast with *permanent virtual circuit*.

virtual computing system—see *virtual machine facility*.

virtual machine facility—an IBM timesharing operating system program that allows multiple remote terminals to access the host.

virtual memory—see *virtual storage*.

virtual mode—in IBM systems, a group of instructions that can be moved in and out of memory.

virtual storage—the amount of space on data storage devices that may be regarded as main storage by the user of a computing system, in which virtual addresses are mapped into real addresses. The size of the virtual storage is limited only by the addressing scheme of the computing system and by the amount of auxiliary storage available, rather than by the actual number of main storage locations.

virtual storage address—a logical memory location.

virtual telecommunications access method—in an IBM 370, a method to give users at remote terminals access to applications programs in a main computer; it also provides resource sharing, a technique for efficiently using a network to reduce transmission costs.

virtual telecommunications access method (VTAM)—an IBM communications I/O control programming software that uses virtual techniques.

visually impaired attendant service—achieved by augmenting the normal visual signals provided on a standard attendant position with special tactile devices and/or audible signals that enable a visually impaired person to operate the position.

VLF—see *very low frequency*.

VLSI—very large scale integration. See *large scale integration*.

voice band—see *voice grade channel*.

voice digitization—conversion of an analog voice into digital symbols for storage or transmission.

voice frequency (VF)—any frequency within that part of the audio frequency range essential for the transmission of speech of commercial quality; i.e., 300 to 3000Hz. Also called *telephone frequency*.

voice grade—telecommunications link with a bandwidth (about 3 KHz) appropriate to an audio telephone line.

voice grade channel—a channel with a frequency range from 300 to 3000 Hz and suitable for the transmission of speech, data, or facsimile.

voice message service—provides the ability for a station user to access an optional voice message recording facility and leave a message for a particular station user.

voice-operated device—a piece of equipment that permits telephone currents to control a task.

voice paging access—allows attendants and station users the ability to have dial access to customer-provided loud-speaker paging equipment.

voice print—a technique for verifying an individual's identity through the pattern produced by his or her voice.

voice recognition—a system of sound sensors that translates the tones of human sounds into computer commands.

voice response unit—a computer-connected device that can selectively link sentences of stored words to create a spoken message.

voice store-and-forward system (VSF)—processor-controlled system that allows voice messages to be created, edited, sent, stored, and forwarded. Users access and operate the system by means of any 12-button dialpad, in response to voice prompts from the system.

voice synthesis—computer-generated sounds that simulate the human voice.

Glossary

volatile storage—memory that loses its contents when electrical power is removed.

volume—a storage medium that can be mounted as a unit; for example, a reel of magnetic tape, a disk pack, or a data cell.

volume table of contents—a table or directory on a direct access storage medium that is mounted as a unit; it describes each data set on the unit.

VRC—see *vertical redundancy check*.

VSF—voice store-and-forward.

VTAM—see *virtual telecommunications access method*.

VTOC—see *volume table of contents*.

W

wait state—the condition of a central processing unit when all operations are suspended.

WARC—World Administrative Radio Conference. ITU conference for deciding the allocation of international radio frequencies and satellite geostationary orbit locations.

warm start—the restart activity appropriate when a temporary failure has not disturbed backup storage.

warning message—an indication during program compilation that a possible error has been detected.

WATS—see *wide area telephone service*.

WATTC—World Administrative Telegraph and Telephone Conference. ITU conference for deciding provisions for the interconnecting and interworking of the world's telecommunications networks.

waveguide—transmission path in which a system of boundaries guides electromagnetic energy. The most common of these are hollow metallic conducting tubes (microwave communications) or rods of dielectric material. See also *fiber-optic waveguides*.

white noise—see *uniform-spectrum random noise*.

wide area telephone service (WATS)—telephone company service allowing reduced costs for certain telephone call arrangements. This can be in-wats, or 800-number service, where calls can be placed to a location from anywhere at no cost to the calling party, or out-wats, where calls are placed out from a central location. Cost is generally based on hourly usage per wats circuit and on distance based on zones, or bands, to which or from which calls are placed.

wide frequency tolerant power plant—provides PBX power facilities that will operate from ac energy sources that are

not as closely regulated as commercial ac power. The wide tolerant plant will tolerate average frequency deviations of up to ± 3 Hz or voltage variations of -15 percent to $+10$ percent as long as both of the conditions do not occur simultaneously. This feature permits operation with customer-provided emergency power generating equipment.

wideband—see *broadband*.

wideband channel—channel wider in bandwidth than a voice grade channel. See also *broadband*.

Winchester disk—a rigid, nonremovable, magnetic oxide-coated, random access storage medium built as a sealed unit.

windowing—see *split screen*.

word—a group of characters capable of being processed simultaneously in the processor; it is treated by the computer circuits as an entity.

word length—the number of bits or bytes that can be processed and stored as a unit.

word processing—the manipulation of text files to create documents.

words per minute—a measure of transmission speed computed on the basis of six characters (five plus a space) per word.

work center—a physical area where a particular type of job is performed.

work file—in sorting, an intermediate file used for temporary storage of information.

working set—the page of a user's program that must be active for the program to operate.

world numbering plan—CCITT numbering plan that divides the world into nine zones. Each zone is allocated a number that forms the first digit of the country code for every country in that zone. The zones are as follows:

- (1) North America, • (2) Africa, • (3 and 4) Europe, • (5) South America, • (6) Australia, • (7) USSR, • (8) North Pacific (Eastern Asia), and • (9) Far East and Middle East.

wpm—see *words per minute*.

wraparound—on a CRT display device, the continuation from the last character position in the display buffer to the first position in the display buffer.

write—to record information on a storage device, a data medium, or an output display.

Glossary

wtnng—waiting.

X

X.21—a technical specification recommended by the CCITT that describes the interface used in the CCITT X.25 packet switching protocol and in certain types of circuit switched data transmissions.

X.25—CCITT recommendation that specifies the interface between user data terminal equipment (DTE) and packet-switching data circuit-terminating equipment (DCE).

X.400—a standard for electronic mail exchange; developed by the CCITT.

X.75—a standard for connecting X.25 networks; developed by the CCITT.

xbar—crossbar.

XD—ex-directory. XD refers to a subscriber number that is not listed in a printed directory. Also called *unlisted number*.

Xenix—a multiuser operating system developed by Microsoft, Inc.; a subset of UNIX; see also *UNIX*.

xerography—a nonchemical photographic process in which light discharges a charged dielectric surface.

xfr—transfer.

xmit—transmit.

xmodem—an eight-bit, public domain error checking protocol developed in the late 1970s by Ward Christensen. The file transfer protocol uses a 128-byte data block and CRC or checksum error checking methods.

X-on/X-off—a method of controlling data communications; it essentially allows a terminal to activate a line when it is ready to receive information and suspend line activity when the terminal is overloaded; see *asynchronous*.

x,y coordinates—the horizontal (row) and vertical (column) designation of a position or dot in an arrangement of elements in perpendicular rows.

x,y recorder—a device that traces on a chart the relationship between two variables, neither of which is time.

Y

yellow pages (YP)—administers files and automatically propagates YP database updates to client systems. This simplifies network administration.

ymodem—a file transfer protocol based on CRC xmodem that was developed by Chuck Forsberg. Ymodem has a 1024-byte packet size.

Z

Z80—an 8-bit microprocessor developed by Zilog, Inc.

zero suppression—the elimination of nonsignificant 0s in a numeral.

zip tone—short burst of dial tone to the headset of an ACD agent, indicating that a call is being connected to the agent console. □

