

Bellcore

 Bell Communications Research

88-76

TECHNICAL ADVISORY
TA-TSY-000885
Issue 1, November 1988

PPSN Support of SNA/SDLC Interfaces

Reproduction of this Technical Advisory is permitted solely for the express purpose of facilitating a unified corporate review. Reproduction of this document for any other purpose is an infringement of Bellcore's copyright and is specifically forbidden.

Copyright ©1988,
Bell Communications Research, Inc.
ALL RIGHTS RESERVED.

TECHNICAL ADVISORY

NOTICE OF DISCLAIMER

This Technical Advisory is published by Bellcore to inform the industry of Bellcore's preliminary view of proposed generic requirements for *PPSN Support of SNA/SDLC Interfaces* and to solicit industry comment thereon. It is a preliminary document subject to review and change. As such, future issues of this information may differ extensively in content and format.

Bellcore reserves the right to revise this document for any reason, including but not limited to, conformity with standards promulgated by various agencies, utilization of advances in the state of the technical arts, or the reflection of changes in the design of any equipment, techniques, or procedures described or referred to herein.

BELLCORE MAKES NO REPRESENTATION OR WARRANTY, EXPRESSED OR IMPLIED, WITH RESPECT TO THE SUFFICIENCY, ACCURACY, OR UTILITY OF ANY INFORMATION OR OPINION CONTAINED HEREIN. BELLCORE EXPRESSLY ADVISES THAT ANY USE OF OR RELIANCE UPON SAID INFORMATION OR OPINION IS AT THE RISK OF THE USER AND THAT BELLCORE SHALL NOT BE LIABLE FOR ANY DAMAGE OR INJURY INCURRED BY ANY PERSON ARISING OUT OF THE SUFFICIENCY, ACCURACY, OR UTILITY OF ANY INFORMATION OR OPINION CONTAINED HEREIN.

This document is not to be construed as a suggestion to any manufacturer to modify or change any of its products, nor does this document represent any commitment by Bellcore, any Bell Operating Company (BOC)* or any regional affiliate thereof to purchase any product whether or not it provides the described characteristics.

Readers are specifically advised that each BOC or regional affiliate thereof may have requirements or specifications different from the generic descriptions herein. Therefore, any vendors, or manufacturers of products should communicate directly with a BOC or regional affiliate thereof to ascertain that company's needs, specifications, and actual requirements.

Nothing contained herein shall be construed as conferring by implication, estoppel or otherwise, any license or right under any patent, whether or not the use of any information herein necessarily employs an invention of any existing or later issued patent.

Bellcore does not recommend products, and nothing contained herein is intended as a recommendation of any product to anyone.

If further information regarding technical content is required, please contact:

Fred Klein, District Manager
Packet Switching Technology
Bell Communications Research, Inc.
331 Newman Springs Road, Room 2X309
Red Bank, New Jersey 07701-7030

For general information, contact:

Information Exchange Management
Bell Communications Research, Inc.
445 South St., Room 2K-122
P.O. Box 1910
Morristown, NJ 07960-1910

* "Bell Operating Company" or "BOC" means a divested Bell Operating Company.

THIS PAGE IS BLANK

PPSN Support of SNA/SDLC Interfaces

CONTENTS

1. Introduction	1-1
1.1 Background on SNA and SDLC	1-1
1.2 Current IBM Support of X.25 and PSPDN Interworking	1-1
1.3 Support of SNA/SDLC and X.25 DTEs	1-2
1.4 Network Capabilities that are Under Consideration	1-2
1.5 Terminology	1-3
2. General Requirements	2-1
2.1 Major Connection Types	2-1
2.2 Options Defining PAD Types: TPAD, HPAD, and SPAD	2-2
2.3 Integrating PAD Functions into the Network	2-3
2.4 Compatibility with SNA/SDLC Products	2-3
2.5 Support of Customer Network Management Tools	2-4
2.6 Billing	2-4
3. PAD SDLC Interface Support	3-1
3.1 Physical Level	3-1
3.2 SDLC Link Layer	3-2
3.2.1 SDLC Protocol and Asynchronous Balanced Mode Extension	3-2
3.2.2 Logical Link Control	3-4
3.2.2.1 Type B Connection Support	3-4
3.2.2.2 QLLC Support	3-5
3.2.2.3 PAD Polling Functions	3-6
4. PAD Virtual Circuit Support	4-1
4.1 Support of PVCs and SVCs	4-1
4.2 X.25 Product Compatibility	4-2
5. Associating SNA/SDLC PUs with Virtual Circuits	5-1
5.1 Mapping LUs & PUs to and from VCs	5-1
5.1.1 General Requirements	5-1
5.1.2 Requirements for Ports with Fixed PU/LU Characteristics	5-3
5.1.3 Requirements for Dial-Up (Switched) Ports	5-4
5.1.4 Requirements for Ports with Grouped PU/LU Characteristics	5-4
5.2 X.121 Address Assignments & Hunt Groups	5-7
5.3 Mapping XID Parameter Values	5-8
5.4 Virtual Circuit Setup/Clearing	5-9
5.4.1 Permanent Virtual Circuits	5-10
5.4.2 Switched Virtual Circuits	5-10
5.4.2.1 All Virtual Calls	5-10

5.4.2.2	PAD-Initiated Calls	5-16
5.4.2.3	DTE-Initiated Calls	5-17
5.5	Virtual Circuit Profiles	5-18
5.6	Data Transfer Phase Functions	5-19
5.7	Security Features	5-19
5.7.1	Access Control	5-20
5.7.2	Disconnection Protection	5-21
6.	SDLC/VC Mappings and State Changes	6-1
6.1	Mapping SDLC Frames to/from Packets	6-1
6.1.1	Default Case: Support of QLLC	6-1
6.1.2	Differences if Type B Connection Supported Without QLLC	6-2
6.2	Send Queues and Flow Control: SDLC and Virtual Circuit	6-2
6.3	PAD Actions in Response to Events & Received Frames/Packets	6-3
6.4	Exception Conditions, State Changes, & Diagnostics	6-4
Appendix A:	Description of the SNA Upper Layers Processor Concept	A-1
Appendix B:	Prototype Cases Used for Performance Specifications*	B-1
Appendix C:	Asynchronous Control Interface Protocol Extension	C-1
Appendix D:	Control Interface Display Formats for Status and Clearing Messages	D-1
Appendix E:	Recommended Formats for Menu-Oriented Command & Display Screens	E-1
References	R-1

LIST OF FIGURES

Figure 1-1. PPSN Support of IBM SNA Products Offering X.25 Interfaces	1-9
Figure 1-2. PPSN Support of Generic SNA Products Offering X.25 Interfaces	1-10
Figure 1-3. PPSN Supports SNA/SDLC Products Directly	1-11
Figure 3-1. SDL Flowchart of SDLC PAD Polling of Secondary Stations	3-10
Figure 5-1. Illustrating SNA Support of Simultaneous Access to Different Hosts	5-22
Figure 5-2. Normal Chronological Relationship Among Physical, Link, and SVC Level Events	5-23
Figure 5-3. Time Relationships Among Link Level and SDLC Establishment Events	5-24
Figure A-1. Illustration of SNA Services Processor Concept for LU Switching	A-3

LIST OF TABLES

Table 1-1. Key SNA Acronyms Used in This Document	1-4
Table 1-2. Descriptions of Selected SNA Terms	1-5
Table 2-1. Protocol Type and PPSN Connection Support Summary	2-5
Table 2-2. Characteristics of Access Interface Types on PPSN	2-7
Table 3-1. Summary of Physical and Link Layers of PAD SDLC Interface Support	3-9
Table 4-1. SDLC PAD Support of SVCs and PVCs	4-4
Table 5-1. Valid and Invalid SDLC PAD Interface Configuration Combinations	5-28
Table 5-2. Triggers for VC Establishment/Clearing and Link Mode Transition Actions	5-29
Table 5-3. Two-Digit (Partial) VC Profiles With a Uniform Definition Across PPSNs	5-31
Table 6-1. QLLC Mapping: From Local SDLC Interface Frames	6-6
Table 6-2. QLLC Mapping: From Virtual Circuit Qualified Data Packets	6-7
Table 6-3. SDLC Frames: Summary of PAD Send Conditions & Actions on Receipt	6-8
Table 6-4. QLLC Packets*: Summary of PAD Send Conditions & Actions on Receipt	6-11

PPSN Support of SNA/SDLC Interfaces

This Technical Advisory (TA) presents Bellcore's preliminary view of proposed generic requirements for support of equipment with IBM Systems Network Architecture (SNA) SDLC interfaces on Public Packet Switched Networks. Ultimately, such requirements are intended to be incorporated into a future issue of the PPSN Generic Requirements (PPSNGR), TR-TSY-000301. All section number references to the PPSNGR or TR-TSY-000301 in this document are based on Issue 1, September 1986. The proposed generic requirements, which are to be used in conjunction with the most recent version of the PPSNGR, are contained in the body of this document.

Abstract

This Technical Advisory presents Bellcore's preliminary view of proposed generic requirements for an SDLC PAD function within the Public Packet Switched Networks (PPSNs) operated by Bell Operating Companies (BOCs). This function would most often be implemented within the Access Concentrator network element, but could also be implemented within a Packet Switch. An SDLC PAD permits end user equipment with an SNA/SDLC interface to communicate with remote equipment via virtual circuits supported by the PPSN. The SDLC PAD converts from the native SNA/SDLC protocol to a protocol that is capable of being carried over a packet network virtual circuit. The version of the PAD described in this document does not perform protocol conversion above the SDLC link layer on the interface to the customer's equipment (i.e., the DTE). Thus, the SDLC PAD is limited to support of virtual circuit connections between two SNA-compatible devices or two devices that are otherwise compatible above the SNA link layer (PPSN connection to DTE via SDLC PAD) and/or the X.25 packet layer (other than SDLC PAD connection to the DTE). [PPSN support of SNA environment X.25 products is also discussed in this document.]

The requirements proposed herein are intended to be ultimately incorporated into a future version of the PPSN Generic Requirements (Bellcore Technical Reference TR-TSY-000301).

Key Words

Generic Requirements	IBM
Packet Assembler/Disassembler (PAD)	Packet Network
PPSN	SDLC
Systems Network Architecture (SNA)	X.25

Technical Highlights of Proposed Requirements

Because of the size and complexity of the requirements presented in the body of this document, the major assumptions and key decisions are summarized in the following list:

1. The focus of this TA is network PAD support of SDLC equipment (the most ambitious of the two levels of network support illustrated in Figures 1-1, 1-2, and 1-3). SDLC PADs are to be provided in PPSN Access Concentrators and may optionally also be provided on Packet Switches.
2. It is assumed that the X.25 interface support already specified in the PPSNGR is adequate to support SNA products that have X.25 interfaces conforming to the 1980/1984 CCITT Recommendation (see sect. 4.2). In the future, user guidelines for configuring PPSN X.25 interfaces to SNA DTEs may be developed. Although not themselves vendor generic

requirements, such "tuning" guidelines could assist users in achieving efficient use of PPSN virtual circuit connections for SNA equipment of IBM and other vendors.

3. Tables 2-1 and 2-2 summarize various types of connections through the PPSN, categorized by the type of equipment protocol at either end. For SNA/SDLC equipment, two major categories of connection are distinguished:
 - Type A: SDLC PAD at one end; either SDLC PAD or SNA X.25 (QLLC) DTE at other end
 - Type B: SDLC PAD at one end; either an X.25 DTE that does not support QLLC or an X.28 PAD at other end
4. For greater flexibility and to accommodate type 2.1 nodes, the TA speaks of an SDLC PAD with different configuration options, rather than of a Host PAD and a Terminal PAD.
5. Requirements are provided for network support of: PVCs; SVCs initiated by the PAD (auto call service subscription option); and SVCs initiated by the customer via an authorized asynchronous control terminal. Auto call service also provides for auto call clearing based on one or more of three time-based criteria. An extension to the X.28/X.29 protocols is specified for SVC establishment/clearing and status message display on asynchronous control terminals. Support of customer-initiated SVCs via commands from the SDLC terminal from which the SVC is to originate is under consideration. (Specific recommendations for support of such a PAD capability may be developed in the future; alternatively, further generic requirements may be judged to be unwarranted).
6. Support of switched SDLC interfaces is currently limited to dial out by the PAD at the secondary station end (as either the X.25 calling or called party). [However, dial in by an asynchronous control terminal used to establish an SVC from/to a dedicated or dial out SDLC device is specified.] Support of dial-in and switched SDLC interfaces at the primary station (i.e., host) end are under consideration by Bellcore for possible future generic requirements.
7. The PAD uses local polling to reduce unnecessary overhead across the virtual circuit.
8. The minimum number of logical channels that should be supported by the PAD per SDLC interface varies with line speed: 4 at 2400bps or less; 16 at 4800-9600bps ; and 32 at speeds above 9600bps. There is also a requirement that the PAD support a number of VCs equal to 100% of the maximum number of secondary stations for which it can be configured.
9. Two new PPSN subscription options are introduced to control acceptance of two X.25 facilities since there is no mechanism to directly notify the SDLC terminal user of their invocation on a per-call basis. These two facilities are Call Redirection Notification and Called Line Address Modified Notification.
10. The requirements specify support of one virtual circuit (VC) per SDLC secondary station (PU). Support of more than one VC per PU (i.e., support of separate VCs for groups of one or more LUs) is categorized as being under consideration because of:
 - the substantial increase in complexity of requirements and implementation associated with support of the higher layers of the SNA protocol stack that would be required
 - possible problems in maintaining transparency of support for customer network management tools, such as IBM's NetView

- IBM SNA support of routing to multiple hosts (see Figure 5-1 and the associated text)

Under consideration as an approach for providing support of multiple VCs per PU (i.e., separate terminals on a cluster controller can simultaneously be associated with different virtual circuits), is use of a user-accessible processor within the network, rather than the SDLC PAD itself, to perform the LU switching function. Such a processor might also be used to provide other higher-level SNA functions, such as asynchronous to 3270 Data Stream protocol conversion.

11. The TA distinguishes three different types of PAD SDLC ports that affect PU/VC mapping and X.121 address assignment:
 - Non-switched, fixed PU/LU characteristics: access by dedicated, real terminals (1 address per PU)
 - Non-switched, grouped PU/LU characteristics: emulation of representative terminal groups for host access by diverse terminals (1 address per SDLC interface or interface hunt group)
 - Switched: dial out at terminal end (point-to-point only, 1 address per interface = 1 address per PU)
12. Following the protocol identifier and other information specified by IBM for the Call User Data (CUD) field of the Call Request packet, the TA specifies 3 added octets for PAD use in assigning an incoming call to the correct grouped PU/LU characteristics port at host end.
13. Support of VC Profiles (partial specifications of the Call Request and Clear Request that are stored in network memory) are introduced to simplify SVC establishment by the customer and reduce PAD memory requirements.
14. Security is provided via CUGs for host access and authorization codes for control interfaces. At the link layer, XID is also supported for switched interfaces. If an SVC is cleared, a new SVC might be established to the same host port, but from a new remote end user, without the host process being aware of the change. To avoid this so-called "session tailgating" problem, a specific PAD disconnection protection mechanism is needed. Details for such a security mechanism are under consideration. Until disconnection protection mechanisms are in place, a customer may choose to restrict non-switched host interfaces to PVCs or to configure interfaces (within the FEP software) as switched, even if the SDLC interface to the PAD is actually non-switched.

1. Introduction

1.1 Background on SNA and SDLC

Systems Network Architecture (SNA) is IBM's architecture to support communications among its data processing products. SNA is defined in terms of generic network components; communication links connecting the components; a number of architecture concepts; and a seven layer protocol structure that is similar to, but distinct from, the Open System Interconnection (OSI) Reference Model. Although not internationally standardized like the OSI Reference Model and the X.25 and X.75 protocols conforming to that model, IBM's SNA is a significant factor in data communications because of the extent of its use and the size of the embedded base of SNA-compatible equipment.

The reader who desires an introduction to SNA is referred to two IBM documents:

- Systems Network Architecture Concepts and Products^[1]
- Systems Network Architecture Technical Overview^[2]

Note that references to SNA and SDLC in this document are intended to apply to the equipment of any vendor, as long as the equipment conforms to the applicable interface and protocol specifications cited in this document. References to specific hardware and software products are included as illustrations, to clarify the generic requirements. Such references do not constitute an endorsement of any type nor do they exclude applicability of the requirements and capabilities described to other compatible equipment.

1.2 Current IBM Support of X.25 and PSPDN Interworking

Although SNA has a sizable market presence, it is not the only significant factor in data communication architectures. This is particularly true as the OSI Reference Model and associated internationally standardized data communications protocols (such as CCITT Recommendation X.25) have reached maturity and grown in use. Thus, successful interworking between the SNA and OSI environments has become important.

As the designer of SNA and a major vendor of associated equipment, IBM's actions concerning SNA/OSI interworking are of interest in this regard. IBM's expansion of the functionality of its SNA products to support OSI (e.g., X.25) interworking reflects the need for such interworking. Among the major network products that IBM has introduced into the SNA environment to support interconnection with X.25 packet switched networks are:

- NCP Packet Switching Interface (NPSI), to permit direct X.25 connections to SNA communication controllers (front end processors). Qualified Logical Link Control (QLLC) is a protocol implemented by NPSI and several other SNA X.25 products. QLLC is carried within X.25 Data packets and permits an X.25 virtual circuit connection to replace a physical SDLC connection within the SNA environment.
- "Upstream" X.25 interface support on the IBM 3710 Network Controller, which provides concentration and/or protocol conversion for attached downstream terminal devices.
- Products with integrated X.25 interfaces or PAD functions, such as IBM's S/36, S/38, and 4700.
- Standalone protocol converters, such as the 5973-L02 Network Interface Adapters (NIAs are no longer actively supported).

1.3 Support of SNA/SDLC and X.25 DTEs

There are two basic levels of support that a public packet network such as a Bell Operating Company (BOC) Public Packet Switched Network (PPSN) could provide for SNA/X.25 interworking. In both cases, the interworking consists of using virtual circuit connections in place of physical SDLC links between SNA-compatible equipment. The two levels are:

1. Offer virtual circuit service via normal DTE/DCE X.25 interfaces to customer equipment that can directly terminate X.25 connections (with or without a separate, customer premises protocol conversion device). If necessary, provide special X.25 configuration options that assure optimal operation, given the characteristics of the X.25 implementation on major SNA environment products. Figure 1-1 illustrates this level of support, using terms applicable to IBM SNA equipment. Figure 1-2 illustrates the same level of support, using more generic terms applicable to SNA environment CPE provided by any vendor.
2. The most ambitious level of support is to also provide protocol conversion support as a network service, so that customer SNA/SDLC devices can directly connect to the packet network without any customer premises protocol conversion equipment. [Note: the term "protocol conversion" is used generically here; the type of SDLC/X.25 interworking function that is the focus of this document could be more accurately characterized as a form of "protocol encapsulation."] This second level of support allows customers to share protocol conversion costs on an economic basis with other users and increases the probability of compatible protocol conversion implementations at the two ends of the virtual circuit connection. Figure 1-3 illustrates the case of network-provided protocol conversion support.

Although there has been growing support for SNA/X.25 interworking on CPE products, economic and operational factors will make direct packet network support of equipment (DTEs) with SNA/SDLC interfaces attractive for many SNA customers. Thus, this document addresses network equipment requirements associated with provision of the second (most complete) level of SNA equipment support. Because the X.25 interface service specified in the current PPSN Generic Requirements (PPSNGR)^[3] addresses the needs associated with the first two levels of support, this document will focus on SNA/X.25 protocol conversion support provided by PPSN PADs, as illustrated in Figure 1-3.

1.4 Network Capabilities that are Under Consideration

Network support of SNA/SDLC equipment specified in this document does not include interpretation or processing of layers of the SNA protocol stack above the SDLC link layer. Such higher layers are passed transparently by the network as user data. Thus, the SDLC PAD is limited to support of virtual circuit connections between two SNA-compatible devices or two devices that are otherwise compatible above the SNA link layer (PPSN connection to DTE via SDLC PAD) and/or the X.25 packet layer (other than SDLC PAD connection to the DTE).

The network requirements for support of SNA/SDLC equipment that are provided in this document do not address the following possible network services:

- SSCP/PU emulation by network PADs, to permit individual LUs of a single SDLC link station address (i.e., a PU) to be handled independently by the PPSN (e.g., assign LUs on the same PU to separate virtual circuits)
- protocol conversion by network PADs that would permit asynchronous terminals to appear as 3270 series display station devices to SNA host applications

- protocol conversion by network PADs that would permit bisynchronous (BSC) devices to successfully communicate with SNA/SDLC devices

Whether generic requirements for network support of such services should be developed and, if so, the details of such requirements are under consideration. If such services are supported, one possibility being considered is to centralize network capabilities associated with SNA protocol layers above SDLC in one or more processors within the BOC network. Such processors would be separate from PADs, but be user-accessible via special DTE address(es), for example. Centralization of such higher layer SNA functions in a limited number of network processors or nodes is being considered because these services are likely to generate varying demand from PAD to PAD, require substantial software complexity, and may have to be modified over time to reflect SNA changes and service improvements.

Appendix A presents a high-level description of how such processors might be used. It is suggestive of a possible approach, to illustrate the concept and contrast it with the alternative of implementing such functions within each SDLC PAD. Comments on the need for these higher layer SNA capabilities and the separate network processor approach broadly outlined in Appendix A are welcome.

1.5 Terminology

The use and interpretation of terms in this document are consistent with the PPSNGR.^[3] Since this is a proposed enhancement to the PPSNGR, it is assumed that the reader is already familiar with the latter document. Terms and acronyms defined or described there (e.g., Access Concentrator or AC) will not be defined again here.

However, this document does introduce a number of concepts and terms that are unique to the SNA environment. Although the reader is assumed to be generally familiar with SNA concepts and terminology, Table 1-1 provides a glossary of acronyms used in this document. Table 1-2 provides brief descriptions of key SNA terms.

Within this document, the phrases *under consideration* and *under consideration for possible future generic requirements* are used interchangeably. They will be used to describe features, capabilities, or technical specifications for which no specific recommendations are presented in the current issue of this document, but which are being considered by Bellcore for possible inclusion in future generic requirements. The expression "for further study" was used in the first issue of the PPSNGR to characterize items that were either subject to further study by CCITT or under consideration by Bellcore for future generic requirements. "Under consideration" is used in this document to clearly distinguish the latter case.

Although specific details are not yet available for "under consideration" items, the label is a useful alerting mechanism, which allows interested parties to:

- provide Bellcore with early technical comments and suggestions that can be used as part of the consideration process
- anticipate possible additional capabilities and/or recommendations for uniformity in their current development efforts (e.g., design software structure or modularity so that it could be more easily extended to accommodate a capability broadly described and identified as under consideration)

Table 1-1. Key SNA Acronyms Used in This Document

ABM	Asynchronous Balanced Mode
ABME	Asynchronous Balanced Mode Extended
ACF	Advanced Communication Function
BNN	Boundary Network Node
BSC	Binary Synchronous Communication (Bisynchronous)
CUD	Call User Data
dm	disconnected mode
DSP	Display System Protocol
ELLC	Enhanced Logical Link Control
FEP	Front End Processor
HDLC	High-level Data Link Control
HPAD	Host PAD
LLC	Logical Link Control
LU	Logical Unit
NCP	Network Control Program
NIA	Network Interface Adapter
NPSI	NCP Packet Switching Interface
NRM	Normal Response Mode
NRME	Normal Response Mode Extended
NRZI	Non Return to Zero Inverted
PSH	Physical Service Header
PU	Physical Unit
QLLC	Qualified Logical Link Control
SDLC	Synchronous Data Link Control
SPAD	Symmetric PAD
SSCP	System Services Control Point
TPAD	Terminal PAD
VTAM	Virtual Telecommunications Access Method

Table 1-2. Descriptions of Selected SNA Terms

1. **Asynchronous Balanced Mode (ABM).** Like SDLC's normal response mode, but for another HDLC subset in which multipoint operation is not permitted, but the stations at both ends of the link simultaneously control the link (primary and secondary roles are not separated). The LAPB link level protocol of X.25 uses this type of balanced operation.
2. **Authorisation Code.** An alphanumeric string of 6 to 10 characters that establishes the identity and authority of a user requesting control interface capabilities on a PPSN SDLC PAD (as specified in this document).
3. **Binary Synchronous Communications (Bisynchronous, Bisynch, or BSC).** A link level character-oriented, synchronous IBM protocol. Has been superseded by the SNA SDLC bit-oriented protocol, but is still widely used because of the large embedded base of BSC equipment.
4. **Boundary Network Node (BNN).** A (type 4 or 5) node which uses network addresses for routing and provides support to attached peripheral nodes.
5. **Call User Data (CUD).** A field of the Call Request packet that contains up to 16 octets (128 octets for Fast Select) of protocol information and/or user data.
6. **Cluster Controller.** A device which controls input and output for multiple connected devices, such as printers and display stations. Within an SNA network, such a device corresponds to an SDLC secondary station address. Also referred to as a *control unit*.
7. **Communication Controller.** See entry for *front end processor*.
8. **Display System Protocol (DSP).** A widely used fourth layer protocol that is carried as user data over an X.25 virtual circuit connection. This protocol permits DTEs whose native protocol is BSC to successfully communicate via an X.25 packet switched connection.
9. **Enhanced Logical Link Control (ELLC).** An LLC protocol, like QLLC, that permits SNA equipment to communicate over a packet network. ELLC is used for peer-to-peer communications (e.g., IBM S/36), and is designed to compensate for possible network failures, such as duplication, resequencing, or loss of packets.
10. **Front End Processor (FEP).** A general term for what is called a *communication controller* within the IBM SNA framework. This is a (type 4) node which performs line control and routing functions within the SNA network, to free the associated host computer of these lower level communications functions so that it may focus on data base and data processing functions. The NCP software runs in these processors.
11. **High-level Data Link Control (HDLC).** An ISO defined link level protocol which includes SDLC and X.25 LAPB as subsets or specific classes of procedure.
12. **Initialization Mode.** See description under *normal response mode*.

Table 1-2. Descriptions of Selected SNA Terms (Cont.)

13. **Logical Link Control (LLC).** A protocol sublayer introduced to provide full IBM SNA link layer functionality across an X.25 packet connection (e.g., to permit XIDs and TEST frames to be supported). There are different types of LLC, depending on the type of IBM equipment involved (e.g., QLLC and ELLC). The type of LLC is identified in the protocol identifier carried as the first octet of the call user data field of the Call Request packet.
14. **Logical Unit (LU).** One of three SNA network addressable units, an LU serves as a port through which an application program or end user can access the SNA network. An LU will typically be associated with a display terminal, printer, or application program process.
15. **NCP Packet Switching Interface (NPSI).** See entry for *X.25 NCP Packet Switching Interface*.
16. **NetView™ *** A host-resident IBM program for network management and problem determination. It consolidates and replaces a number of previous IBM software products, such as NPDA.
17. **Network Addressable Unit (NAU).** A source or destination of information being passed through an SNA network. There are three types of NAU: LU, PU, and SSCP.
18. **Network Control Program (NCP).** An IBM program that runs in a FEP and provides communication controller support.
19. **Network Interface Adapter (NIA).** An IBM protocol conversion device which employs the PSH protocol to permit SNA equipment to communicate over an X.25 packet network. This device is no longer actively supported by IBM.
20. **Normal Disconnected Mode (NDM).** See description under *normal response mode*.
21. **Normal Response Mode (NRM).** A secondary station on an SDLC link may be in one of three modes: initialization mode, normal response mode, or normal disconnected mode. The initialization mode is for (re)initializing the station, using product specific procedures. This mode is not supported by the requirements of this document. The normal response mode is entered in response to receiving a SNRM command from the primary. It is the mode during which data transmission is conducted. The normal disconnected mode is entered in response to receiving a DISC command from the primary and is the initial state after power on.
22. **Physical Services Header (PSH).** A protocol used by IBM Network Interface Adapters to encapsulate SNA communications within X.25 packets so that they can be transported across an X.25 packet network. PSH corresponds to one of the LLCs defined by IBM.

* NetView is a trademark of the International Business Machines Corporation.

Table 1-2. Descriptions of Selected SNA Terms (Cont.)

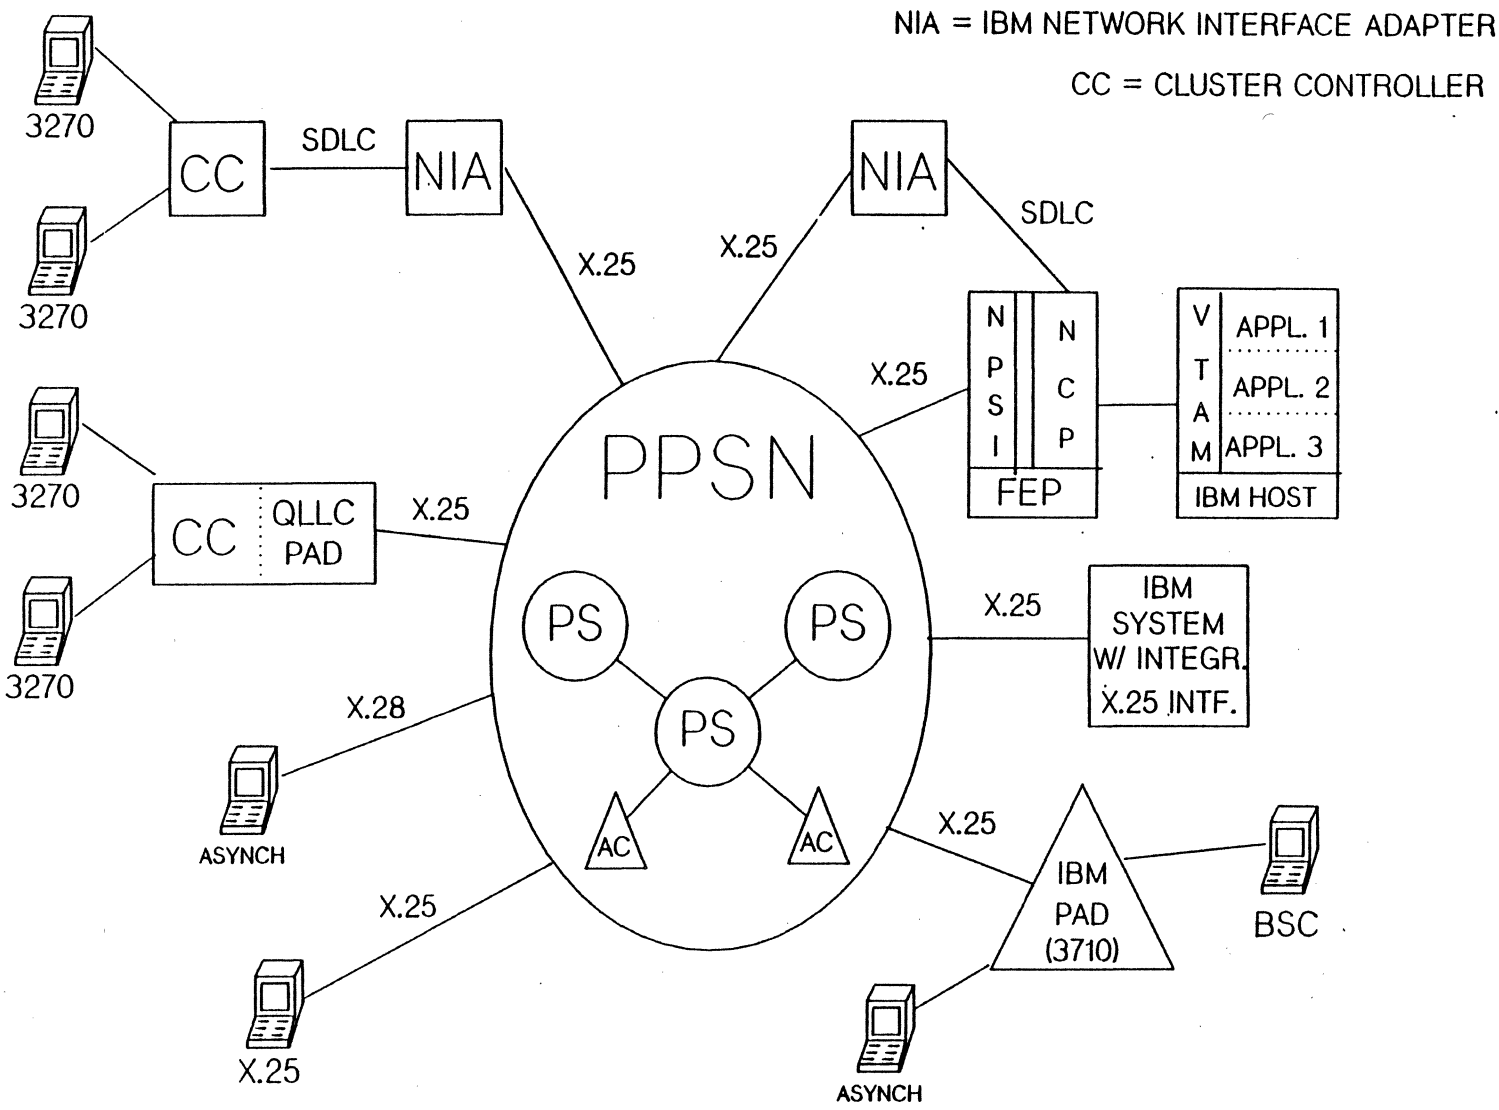
23. **Physical Unit (PU).** One of three types of SNA network addressable units, a PU manages and monitors the resources of a node. Although the PU is really a set of services rather than a piece of hardware, there is typically a one-to-one correspondence between the PU functions and an SDLC station address device, such as a cluster controller.
24. **Polling.** In the SDLC protocol, secondary stations transmit only one at a time and only when explicitly authorized by the primary station on the link. The primary station authorizes responses from a secondary station by polling that station. The primary station polls by setting the P (poll) bit in a command frame addressed to the designated secondary station.
25. **Primary.** A master/slave relationship applies to SDLC links. One end of the link is the master or *primary* station, which is responsible for controlling the link and data flow over the link in both directions. One or more stations at the other end of the SDLC link serve in a secondary capacity. Only one secondary station may transmit at any given time, and only in response to polling commands from the primary station. Since there can be only one primary station per SDLC interface, link station addresses always identify the secondary station that is being addressed or that is responding.

The terms primary and secondary can also be applied to PUs and LUs (resident at a link station address).
26. **Protocol Identifier.** For virtual circuit connections involving IBM SNA equipment, the first octet of the call user data field of the Call Request packet is used to identify the type of connection and LLC protocol being used. This octet is called the protocol identifier.
27. **Qualified Logical Link Control (QLLC).** A protocol sublayer carried within qualified (i.e., Q bit set to 1) X.25 Data packets. QLLC allows a packet network virtual circuit connection to replace a direct SDLC connection without loss of SDLC functionality that is not directly supported by X.25 (e.g., XID). QLLC is one of the LLCs defined by IBM. It is employed by NPSI and is key in IBM's support of X.25 interoperability.
28. **Secondary.** See the description of *primary*.
29. **Synchronous Data Link Control (SDLC).** The SNA link level protocol. SDLC is a synchronous, bit-oriented protocol, which is the successor to BSC. It is a subset of HDLC that supports both point-to-point and multipoint operation, with distinct primary and secondary station roles.

Table 1-2. Descriptions of Selected SNA Terms (Cont.)

30. **System Services Control Point (SSCP).** One of three types of network addressable units, the SSCP is the central point within the SNA network for configuration management, session establishment, and network operations and maintenance coordination.
31. **Virtual Circuit (VC) Profile.** A partial specification for virtual call establishment and clearing that is stored by the PPSN and can be referenced by number. This concept is introduced in this document to provide a mechanism to facilitate call establishment and reduce memory requirements for storing configured call data.
32. **Virtual Telecommunication Access Method (VTAM).** An IBM host-resident program that provides a unified SNA communication network interface for multiple applications within the host. Other access methods are in use (e.g., TCAM), but VTAM is the one which is expected to have full IBM support into the future.
33. **X.25 NCP Packet Switching Interface (NPSI).** IBM software resident in a FEP which permits that FEP to directly support X.25 interfaces.

Figure 1-1. PPSN Support of IBM SNA Products Offering X.25 Interfaces



2. General Requirements

2.1 Major Connection Types

There are two basic types of connections through a packet network involving an SNA SDLC DTE at one end and another SNA SDLC DTE or an X.25 DTE at the other end:

- **Type A.** In this connection, there is an SNA/SDLC DTE at one end of the connection as seen by the PPSN and either another SNA/SDLC DTE or an X.25 DTE supporting IBM QLLC at the other end. An SDLC PAD is required at each end that has an SNA/SDLC DTE, to convert between SDLC and the native packet protocol of the network node for virtual circuit service across the network. Each PAD communicates with the remote PAD or X.25 DTE using the IBM QLLC logical link control protocol that is passed transparently by the packet network as qualified user Data packets. This communication creates a logical link level connection between the two DTEs, in spite of the intervening packet network(s).
- **Type B.** There is an SNA/SDLC DTE at one end of the connection and either a native X.25 DTE (as seen by the PPSN) that does not support QLLC or an asynchronous DTE attached via an X.28 interface to PPSN at the other end. An SDLC PAD is required only at one end for protocol conversion. Unlike the Type A connection case, the PAD does not use a fourth layer protocol to communicate with the remote end. The SDLC link is terminated at the local SDLC PAD, which protocol converts the SDLC frames to virtual circuit packets for transmission to the remote X.25 or X.28 DTE. Although the SDLC PAD provides SDLC/packet protocol conversion, the two DTEs must nevertheless be able to communicate successfully with one another at protocol layers above the SDLC and X.25 packet layers, respectively.

The mode of operation of an SDLC PAD is determined by the type of connection and the type (primary, secondary, or peer) of DTE to which the PAD is attached. A PAD in a type A connection supports SDLC towards the local DTE and QLLC towards the remote end. A PAD in a type B connection supports SDLC towards the local DTE and an X.25-compatible virtual circuit toward the remote end DTE. Thus, the nature of the protocol conversion performed by the PAD varies with the connection type. The difference in PAD function associated with the type of SDLC device to which it is attached are detailed in section 2.2.

Table 2-1 summarizes the various types of connections possible across a PPSN, categorized by type of access interface on either end of the connection. For each connection, the protocol(s) that are passed within the network(s) are indicated, if the connection is supported by PPSN. The first two interface columns and first two interface rows of the table cover the two connection types identified above for those cases in which an SDLC PAD is used. Table 2-2 summarizes the characteristics of each access interface type identified in Table 2-1.

As indicated by the entries of Table 2-1 other than those of the first two columns or rows, other types of connections involving SNA or SNA-compatible DTEs are possible across a packet network. However, such connections are either handled by the network in accordance with the X.25, asynchronous, or 3270 BSC DSP interface specifications of PPSNGR sections 3.2, 3.3, or 3.8, respectively, or are not directly supported on PPSN. [Note: Section 3.8 does not exist in Issue 1 of the PPSNGR, but proposed text for this new section is presented in Issue 2 of TA-TSY-000301,^[4] section 6.] When an IBM protocol is involved in such cases (e.g., ELLC or PSH), it is carried transparently by the network between the two DTEs.

SNA equipment and software that support connections to packet switched networks distinguish the type of connection to which a virtual call corresponds by means of a protocol identifier

embedded in the call user data (CUD) of the Call Request packet (see sections 3.2.2 and 4.2). In some cases, an SDLC PAD may append details on local PU and LU types to the CUD (see section 5.4.2.1).

2.2 Options Defining PAD Types: TPAD, HPAD, and SPAD

Generically, the term "PAD" is used to refer to the network functionality that performs protocol conversion and packet assembly/disassembly to permit a non-X.25 DTE to communicate via the PPSN. Typically, for asymmetric protocols such as BSC or SDLC, which have a master/slave or primary/secondary distinction, separate PAD functions are identified. The Terminal PAD (TPAD) connects to devices which act as secondary (slave) stations, such as cluster controllers and terminal devices. The Host PAD (HPAD) connects to equipment which act as primary (master) stations, such as mainframe computer front end processors (FEPs). This distinction has been useful because of the differences in virtual circuit establishment, polling, flow control, and frame handling procedures that apply to PADs servicing the two link level station roles, respectively.

However, recent extensions to IBM SNA allow for peer-to-peer communications (i.e., node type 2.1 to node type 2.1) in which primary and secondary roles can be negotiated. In addition, the IBM SNA link level protocol has been extended to accommodate the asynchronous balanced mode (ABM) for balanced stations, as well as the normal response mode (NRM) of traditional SDLC. Because of these extensions to SNA and the fact that TPAD and HPAD distinctions can be introduced as configuration alternatives on a single PAD implementation, a more flexible approach to PAD distinctions is adopted in this document. It is assumed that there is a single PAD function with configuration options that are appropriately set to correspond with the type of device to which it is connected. For some device types, other PAD parameters may be dynamically altered, based on the type of device on the remote end of the virtual circuit connection.

Depending on the type of device to which a PAD SDLC port is connected, the port may be configured for connection to a primary station, secondary station, or negotiable station (for which the primary/secondary role may be negotiated). Also depending on the capability of the attached equipment, the PAD port may be configured for NRM only or NRM/ABM-capable. Finally, the PAD would have configuration options controlling the means available for establishing virtual circuits from or to the SDLC port. Options would indicate whether or not each of the following were supported:

- Permanent Virtual Circuits (PVC)
- Switched Virtual Circuit (SVC) automatic call service
- SVC establishment via asynchronous control terminal
- SVC establishment via in-band SDLC terminal messages (under consideration)

In terms of these PAD options, the traditional SDLC TPAD would correspond to secondary station DTE, NRM only, with support selected from among PVC, automatic call service, and/or per-call SVC establishment. A traditional HPAD would correspond to primary station, NRM only, with support of PVCs and/or automatic call service. A PAD configuration capable of supporting IBM SNA node type 2.1 devices for peer-to-peer connections would be combined station, NRM/ABM-capable, with support of PVCs, automatic call service, and/or some of the per-call SVC options. Such a PAD might be labeled a symmetric PAD (SPAD) for convenience.

2.3 Integrating PAD Functions into the Network

1. Within the constraints of overall resource limitations, an AC should be capable of being provisioned with the number and mix of SDLC PAD ports and other PAD (e.g., asynchronous, 3270 BSC DSP) and X.25 concentration ports that reflect Bell Operating Company (BOC) customer demand at the location served by the AC.
2. A supplier may optionally support SDLC PAD ports within PSs, in addition.
3. Common AC and PS functions, such as routing and OA&M should cover any SDLC port provisioned on that equipment, as specified in the PPSNGR. In particular, any SDLC PAD implementation should comply with the Operations Technology Generic Requirements^[6] (OTGR) as they apply to the PPSN network element containing the SDLC PAD function.
4. An SDLC PAD port supports an SDLC interface in the direction of the DTE and a virtual circuit connection consistent with X.25 (as specified in the 1984 CCITT Recommendation and section 3.2 of the PPSNGR) in the direction of the network. As specified later, the QLLC protocol may be carried over this virtual circuit connection. The actual inter-nodal and inter-network protocol(s) supporting the X.25 virtual circuit connection between two PADs or between a PAD and an X.25 DTE depend on the nature of the network interfaces involved. Within the network, interfaces between nodes of the same vendor may use a proprietary protocol, as long as this protocol can support the X.25 virtual circuit service on an end-to-end basis. Interfaces between nodes of different manufactures should be X.75' within the BOC network and X.75 between networks.

2.4 Compatibility with SNA/SDLC Products

The SDLC PAD should be compatible with SDLC interfaces on all type 1 and 2.0 SNA nodes at the secondary station end and all type 4 SNA nodes at the primary station end. For interfaces configured to support peer-to-peer communication (with negotiable primary/secondary role) the PAD should be compatible with SDLC interfaces on all type 2.1 SNA nodes. Specifically, the SDLC PAD should be compatible with SDLC interfaces on the IBM model 3705, 3720, 3725, and 3745 communication controllers; 3270 Information Display System cluster controllers (i.e., models 3174, 3271, 3274, 3275, and 3276); and the IBM equipment listed below:

Series/1	System/34
System/36	System/38
System/88	3601, 3602
3614, 3624, 2631, 3632	3651, 3661, 3684, 3694
3767	3771, 3774, 3775, 3776
3791	4701, 4702, 4730
5285	5520
5550	8130, 8140, 8150
8775	

The SDLC PAD should also be compatible with products of other vendors that conform to the SDLC interface specifications referenced in this document.

2.5 Support of Customer Network Management Tools

IBM and other vendors provide network management tools that can be used by customers to monitor and control their equipment and associated private networks. For the SNA environment, an example of such a tool is IBM's network management software, NetView, which provides the following types of services for the customer: network administration, data gathering and analysis, fault detection, line monitoring, performance monitoring and rerouting. It is important that the operation of the SDLC PAD does not interfere with the transparent operation of such network management tools.

When an SDLC interface is replaced by a virtual circuit (VC) through a PPSN using an SDLC PAD at one or both ends, the access interface to the DTE should continue to conform to SDLC specifications and the VC should transparently pass all communications of NetView and other network management tools that are carried in the information field of SDLC frames. Therefore, control and monitoring of SNA devices should be supported transparently even when PPSN VCs with SDLC PADs replace one or more SDLC interfaces between these devices. However, tools that rely on special signaling schemes over the SDLC facility will not be supported when a VC replaces a dedicated facility. An example of such a support exception is an SDLC interface with IBM advanced function modems (AFM's) that pass NetView-related information using specialized inband analog signaling. In this case, such analog signals cannot pass through the PPSN and will terminate on the modem at the near (rather than the far) side of the VC.

1. The SDLC PAD should not interfere with the function of customer network management tools, to the extent that they function by passively monitoring the SDLC access interface and/or transmitting messages within SDLC frame information fields.

2.6 Billing

Billing requirements in section 8 and Appendix A of the PPSNGR apply to SNA/SDLC interfaces supported on the PPSN. In particular, the value "15" for BCD characters 2-3 in Tables 178 and 179 of Appendix A in the PPSNGR identify SDLC as the access interface service type.

Table 2-1. Protocol Type and PPSN Connection Support Summary

Orig. End Type	Interface Type at Destination End							
	SDLC PAD, Type A Conn.	SDLC PAD, Type B Conn. (D)	Native X.25	Asynch.	NPSI (or X.25 w/QLLC)	ELLC over X.25	NIA PSH	BSC DSP PAD
SDLC PAD Supporting QLLC (Type A)	QLLC+	NS	NS	NS	QLLC+	NS	NS	NS
SDLC PAD Supporting Simple X.25 (Type B) (D)	NS	A	X.25	X.25, X.29(D)	A	NS	NS	NS
Native X.25 (No QLLC or ELLC)	NS	X.25	X.25	X.25, X.29	X.25+	NS	NS	NS
Asynch. PAD	NS	X.25, X.29(D)	X.25, X.29	X.25, X.29	X.25+, X.29	NS	NS	NS
IBM NPSI (or X.25 with QLLC)	QLLC+	A	X.25+	X.25+, X.29	QLLC	ELLC	PSH	Note 1
IBM ELLC Over X.25	NS	NS	NS	NS	ELLC	ELLC	NS	NS
IBM NIA PSH	NS	NS	NS	NS	PSH	NS	PSH	NS
3270 BSC DSP PAD	NS	NS	NS	NS	Note 1	NS	NS	DSP

(Legend on next page)

Table 2-1. Protocol Type and PPSN Connection Support Summary (Cont)

Legend

NS	Not supported on PPSN
A	Type A connection should be used instead in this case
D	PPSN support is desirable
X.25+	X.25 with SNA protocol ID in Call User Data (CUD)
QLLC+	QLLC, possibly supplemented with PU/LU details in CUD (sect. 5.4.2.1)
Note 1	NPSI Connection to a DSP PAD Would Require Special User Software

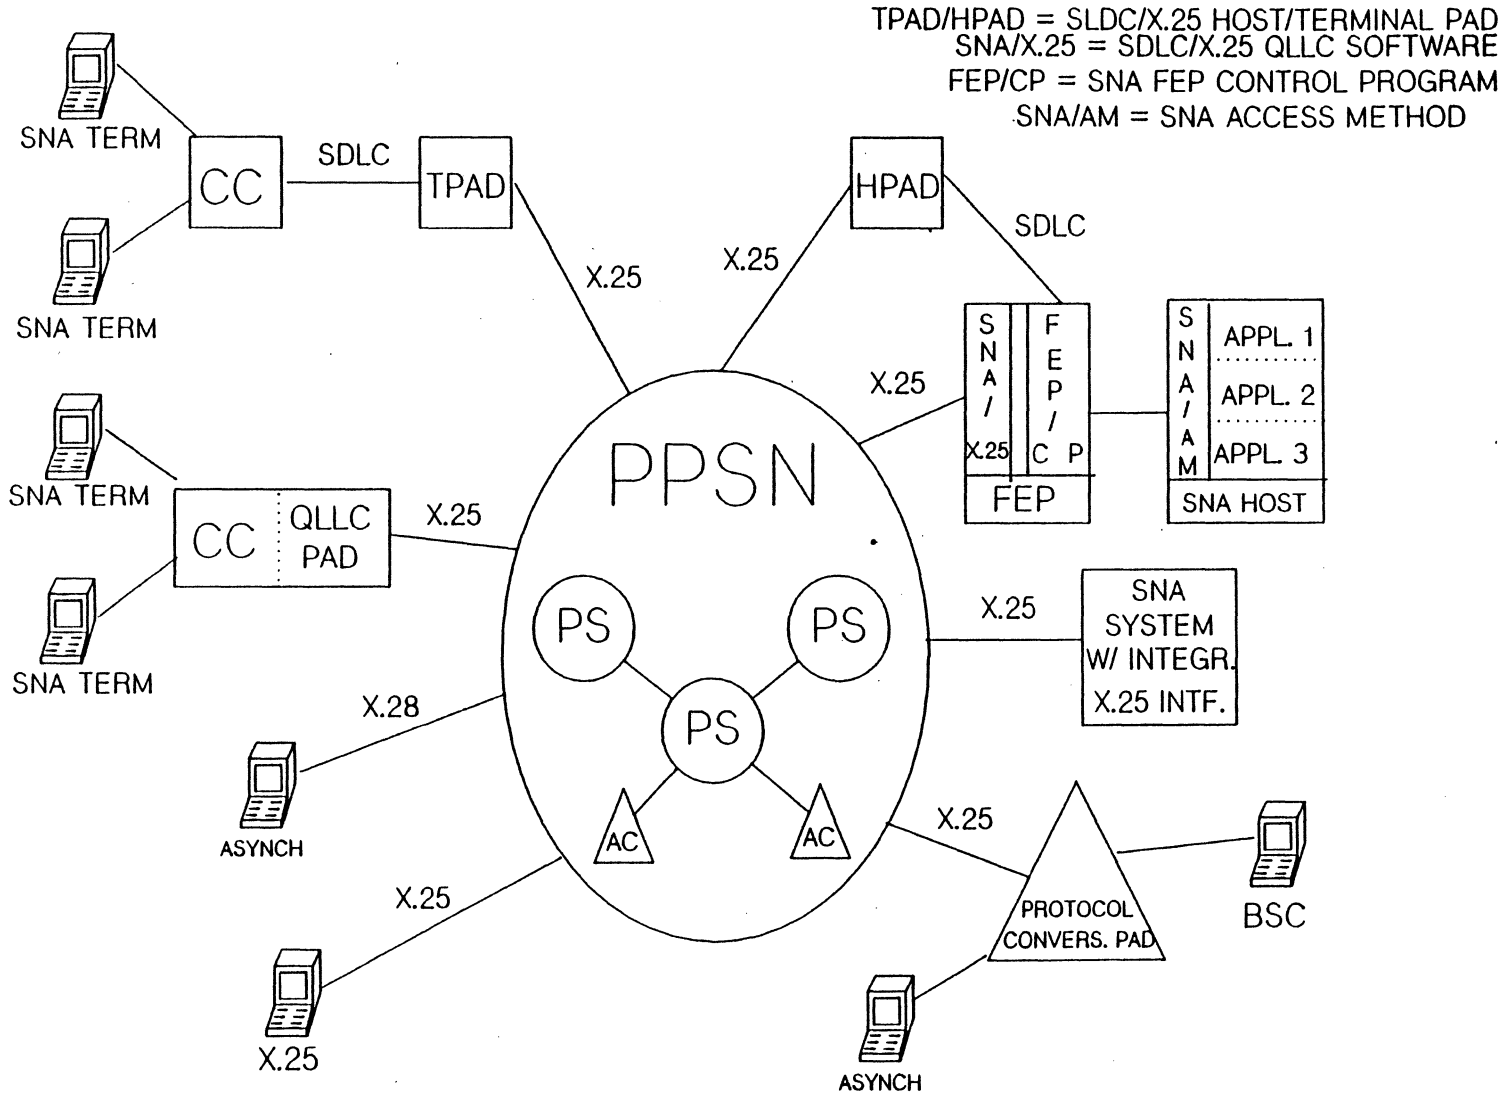


Figure 1-2. PPSN Support of Generic SNA Products Offering X.25 Interfaces

CC = CLUSTER CONTROLLER
 FEP = FRONT END PROCESSOR

SNA/X.25 = SDLC/X.25 QLLC SOFTWARE
 FEP/CP = SNA FEP CONTROL PROGRAM

SNA/AM = SNA ACCESS METHOD

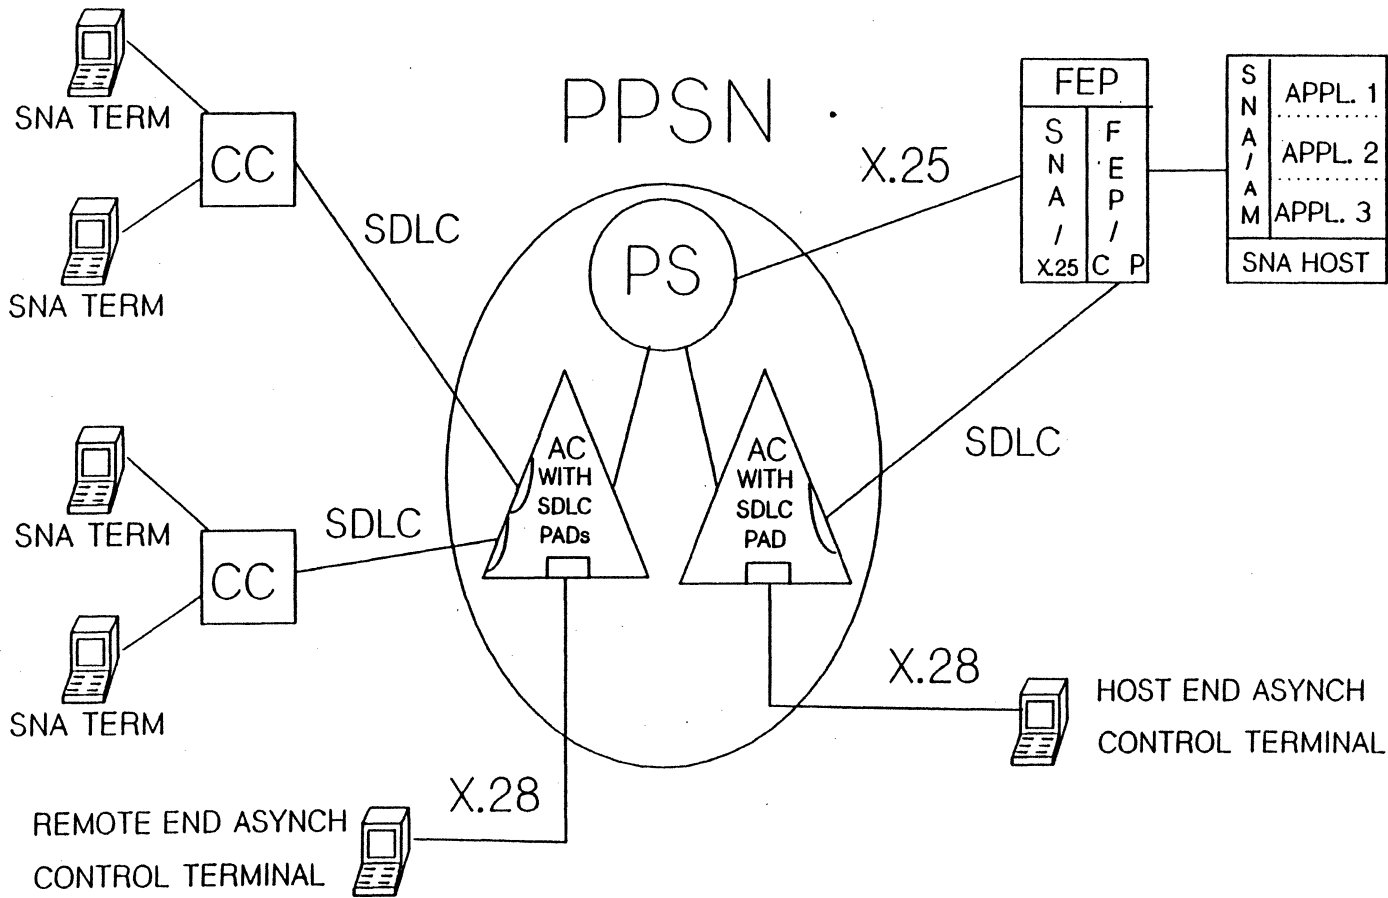


Figure 1-3. PPSN Supports SNA/SDLC Products Directly

Table 2-2. Characteristics of Access Interface Types on PPSN

Access Type	PPSNGR Sect.	PPSN Support	SNA Category	QLLC-LLC Type
SDLC PAD Supporting QLLC (Type A Connection)	3.9*	Req.	SNA	LLC 3
SDLC PAD Supporting Simple X.25 (Type B Connection)	3.9*	Des.	Non-SNA	LLC 0
Native X.25 (No QLLC or ELLC)	3.2	Req.	Non-SNA	LLC 0
Asynch. X.3/X.28/X.29 PAD	3.3	Req.	Non-SNA	LLC 5
IBM NPSI (or X.25 with QLLC)	3.9*, 3.2	Transp.	SNA	All (or LLC 3)
IBM ELLC Over X.25 (e.g., S/36, S/38)	NA	Transp. (need mod. 128)	SNA	LLC 6
IBM NIA PSH	NA	Transp.	SNA	LLC 2
3270 BSC DSP PAD	3.8**	Req.	Non-SNA	LLC 4***

Legend

- Transp. Carried transparently by PPSN
- Req. Support required on PPSN
- Des. Support desirable on PPSN
- NA Not Applicable
- * Assumes this document will be incorporated into future issue of PPSNGR as section 3.9.
- ** BSC DSP requirements are expected to be incorporated into Issue 2 of the PPSNGR. (See sect. 6 of Bellcore TA-TSY-000301, Iss. 2).
- *** It may be possible to use the IBM GATE function (associated with LLC 4), in conjunction with custom SNA host software to communicate with a PPSN BSC PAD implementing the DSP protocol.

BLANK PAGE

3. PAD SDLC Interface Support

Table 3-1 summarizes the characteristics of the physical and link levels of the local SDLC interface supported by the PAD. The details of this support are provided in the following two sections.

3.1 Physical Level

1. All SDLC PADs should support non-switched (i.e., "dedicated") interfaces.
2. Support of dial-out interfaces should be configurable on ports optioned for secondary or negotiable role SDLC stations, that are also optioned for any type of SVC service (including automatic call service). Configuration options for dial-out interfaces should be settable by service order or specifications provided via authorized customer control terminals (see section 5.4.2.3).
3. Support of dial-in interfaces may be needed for ports optioned for per-call establishment of SVCs, in-band over the SDLC interface. Requirements for SDLC dial-in and support of in-band SVC establishment over an SDLC interface are under consideration.
4. Line speeds of 300, 1200, 2400, 4800, 9600, 19200, and 56000 bps should be supported on dedicated interfaces. Line speeds of 300, 1200, 2400, 4800, and 9600 should be supported on switched interfaces. Support of other line speeds, in addition, is desirable.
5. The interface should conform to EIA 232D (CCITT V.24) for line speeds of 19200 bps or less. The interface should conform to CCITT V.35 for line speeds greater than 19200 bps. Optionally, other interface types (such as CCITT X.21) may also be supported.
6. The modem interface and access should conform to Section 3.1.1 and 3.3.2 of the PPSNGR.
7. For 4800 bits/sec and 9600 bits/sec dial-out interfaces, the PAD's modem should be EIA 232D compatible and conform to CCITT Recommendation V.32.
8. For switched interfaces, the PAD should disconnect the established call connection through the PSTN if the received carrier is interrupted or lost for more than 2 seconds. The PAD should also clear any virtual circuit(s), and disconnect and deactivate the link layer associated with the dial connection. Until the latter actions are taken, the PAD should prevent any new dial-up connections from being established through the same port. Any subsequent dial connections should require the same link initialization procedures over the DTE/PAD interface (e.g., XID) normally required for new physical connections.
9. If the received carrier is interrupted or lost for less than 0.6 second, the PAD should not disconnect an established call connection through the PSTN, unless the associated dial-up port is explicitly configured as "no call waiting," in which case the interruption limit should be a BOC-settable value in the range 0.1 to 0.5 second, in 0.1 second increments.
10. Both duplex and half-duplex operation should be supported as interface configuration options.

3.2 SDLC Link Layer

The link layer of the interface supported by the SDLC PAD toward the DTE (the SNA or SNA-compatible customer equipment) is SDLC. For the basic operation of the PAD, protocol layers above the SDLC link layer are transparent to the network PAD on the interface toward the DTE. Support of extended (non-basic) features may require that the network PAD handle higher layers of the SNA protocol across the DTE/DCE interface. Requirements for extended SDLC PAD features, such as support of multiple virtual circuits per PU (SNA Physical Unit), are under consideration.

3.2.1 SDLC Protocol and Asynchronous Balanced Mode Extension

If the SDLC PAD is configured for normal response mode (NRM), then SDLC, in conformance with the IBM document on IBM SDLC Concepts^[6] is supported across the DTE/DCE interface.

In support of communication between SNA nodes of type 2.1 (i.e., peer-to-peer communication), extensions of SDLC that conform to HDLC are defined. These extensions are primarily in support of balanced stations that are capable of negotiating a role as either a primary or secondary station. The exchange identification (XID and XID3) mechanism is used to indicate and negotiate these roles. Further details on type 2.1 nodes and the associated procedural details are available in IBM documentation.^[7]

Specific HDLC extensions to the basic SDLC protocol are as follows:

- Support of the asynchronous balanced mode (ABM), in addition to the normal response mode (NRM).
- Support of the SABM and SABME frame formats, in addition to the SNRM and SNRME frame formats. These additional frame types are used to request balanced station operation (ABM) for modulo 8 and modulo 128 cases, respectively.

Procedures for asynchronous balanced mode operation (ABM and ABME) are not explicitly addressed in IBM's basic SDLC reference document. These procedures are addressed in a separate IBM document.^[8]

1. Link level procedures and all frame formats specified in the IBM documents on SDLC^[6] and its extension to ABM mode^[8] are to be supported by the PAD, unless otherwise stated below. Both normal response mode (NRM) and asynchronous balanced mode (ABM) should be supported by the SDLC PAD as interface configuration options. For dial-up interfaces, the setting of this option should be possible via the XID procedures detailed later in this document.
2. The SDLC PAD should support SNRM, SNRME, SABM, or SABME, depending on whether the interface is configured for NRM modulo 8, NRM modulo 128, ABM modulo 8, or ABM modulo 128, respectively.
3. The SDLC PAD should not generate a REJ frame, but should be able to correctly respond to a valid REJ frame received over the local SDLC interface.
4. Point-to-point operation should be supported for switched interfaces. Both point-to-point and multipoint operation should be supported for dedicated interfaces. In multipoint operation, the SDLC PAD should be able to accommodate up to 4 stations for line speeds of 2400 bps or lower, up to 16 stations for line speeds of 4800-9600 bps and up to 32 stations for line speeds above 9600 bps. For multipoint operation, simultaneous support of a larger number of stations, for each stated line speed, is desirable.

5. Two-way alternate operation and two-way simultaneous operation should be supported as interface configuration options. On dial-up interfaces, the selection between the two options should be possible via the XID procedures detailed later in this document.
6. The SDLC PAD should support connection to a primary station, secondary station, or negotiable role station as an interface configuration option. For dial-up interfaces, the setting of this option should be possible via the XID procedures detailed later.
7. Each PAD SDLC interface should be configurable as NRM only or NRM/ABM-capable. For dial-up interfaces, this option should be settable through the XID procedures detailed later.
8. The maximum frame size supported on the interface should be settable at least to values in the ranges that would accommodate information format frames with information fields of 128-137 octets and 256-265 octets. Support of values in the range that would accommodate information format frames with information fields of 512-521 octets, in addition, is desirable. For switched interfaces, it should be possible to set this value through the XID procedures detailed later.
9. The SDLC PAD should support modulo 8 link level sequence numbers. It is desirable that the SDLC PAD support modulo 128 link level sequence numbers as an interface configuration option. If the interface is configured for modulo 8 link level sequence numbers, link level window values in the range 1-7 should be supported. If the interface is configured for modulo 128 link level sequence numbers, link level window values in the range 1-60 should be supported. Support of window values in the range 61-127, in addition, is desirable for modulo 128 operation.
10. It should be possible to designate separate link level acknowledgment window values in the range supported as a configuration option for each secondary station on the interface. For switched interfaces, it should be possible to set this value via the XID procedures detailed later.
11. The SDLC PAD should retain copies of frames that may need to be re-transmitted, until their receipt is acknowledged.
12. The following SDLC features need not be supported by the PAD:
 - The secondary station initialization mode, and the associated RIM and SIM frame formats.
 - Loop configuration operation, and the associated UP, CFGR, and BCN frame formats.
 - The UI (unnumbered information) frame format, which is not accommodated by IBM's QLLC protocol (used between an SDLC PAD and either an X.25 DTE or another SDLC PAD over a virtual circuit connection)
 - Invert-on-zero transmission coding (also known as Non-Return to Zero Inverted, NRZI).
 - Group addresses, which identify groups of more than one station per interface (i.e., an address which is associated with more than one link station). However, the broadcast address, to which all stations on an interface may respond, should be supported (see item 3 of section 5.1.1).

13. Refer to sections 6.1 and 6.3 for specification of the conditions under which the PAD transmits specific frames and PAD action in response to the receipt of specific frames.

3.2.2 Logical Link Control

Logical link control refers to the protocol used to support a virtual link between the end points of a virtual circuit, despite the presence of one or more intervening packet networks. This protocol rides above the X.25 packet layer and permits certain SDLC link level commands and responses to be passed between the DTEs, as if there were a direct physical connection between them. PPSN SDLC PAD support is defined for two logical link control options, corresponding to the two connection types identified earlier:

- **Type A:** support of QLLC (with or without PU/LU details added to the call user data of the Call Request Packet, following the protocol identifier)
- **Type B:** support of no logical link control protocol by the network PAD

Which of these two cases apply for a given virtual circuit is specified as an interface configuration option for PVCs and is specified by the protocol identifier present in the first octet of the call user data (CUD) field of the Call Request packet for SVCs. The protocol identifier specified by an IBM NPSI X.25 interface or other X.25 interface initiating a packet call is assumed to correctly reflect the type of device on the remote end. For an interface which is the destination for an SVC, the SDLC PAD uses the protocol identifier in the CUD of the received Call Request packet. Unless a type B connection CUD value is specified, as the interface configuration option or via control terminal command for the SVC, the SDLC PAD assumes a type A connection and inserts a protocol identifier corresponding to QLLC (LLC 3) in any SVC it initiates. (A first octet value of X'C3' indicates QLLC; all other values are other than QLLC). Thus, the SDLC PAD cannot initiate an SVC to a DTE that does not support QLLC (e.g., a native X.25 or X.28 DTE), unless the calling party is aware, in advance, that the connection will be type B.

If QLLC is supported over the virtual circuit, certain frames passed over the local SDLC interface are mapped from or to corresponding QLLC commands or responses carried to the remote end in qualified Data packets. Rather than reacting to these frames locally, the SDLC PAD awaits QLLC packets from the remote end, which are then mapped to corresponding frames on the local SDLC interface. Details of QLLC support are presented in sections 3.2.2.2, 6.1.1 and 6.3.

For Type B connection types, QLLC is not supported by the SDLC PAD over the virtual circuit. For Type B connections, all SDLC frames are handled locally, with the exception of the user data in I-frames, which is mapped to Data packets over the virtual circuit connection. Additional details for this case are provided in the next subsection.

In both of the above cases, polling functions are handled locally by the PAD and are not reflected across the virtual circuit connection. This reduces the overhead of the virtual circuit and improves performance. Details of how the SDLC PAD handles polling are provided in section 3.2.2.3.

3.2.2.1 Type B Connection Support

Type B connection support refers to the SDLC PAD capability to support a virtual circuit connection with the remote end over which QLLC is not used. Under this option, the PAD handles the SDLC protocol exclusively on a local basis. This is not to say that there is no propagation of user data or flow control between the local SDLC link and the virtual circuit connection to the remote end. However, in the absence of QLLC, the normal degree of separability in handling adjacent links at layer two (e.g., LAPB of X.25) applies to the local SDLC link and the

virtual circuit connection to the remote end. [In other words, acknowledgments at the link layer are local to that single link, rather than performed on an end-to-end basis.] In particular, responses to SDLC commands across the local link are not delayed until corresponding responses from the remote end of the virtual circuit are received (as is done via the QLLC mechanism in the case of Type A connections).

1. Support of Type B connection support option on a PAD interface is desirable. If supported, the following apply.
2. Selection of this option is by interface configuration for PVCs. For SVCs, support of an interface-specific default assumption (QLLC or Type B connection support), which can override the network default assumption of QLLC, is desirable. In either case, if the first octet of the protocol identifier in the CUD of the Call Request packet is set to any hexadecimal value other than "C3" then the Type B connection support option is assumed for that SVC.
3. The PAD interface to the DTE conforms to section 3.2.1.
4. Handling of XID frames across the local SDLC interface, conforms to section 5.3, with the further constraint that QLLC is not supported.
5. The SDLC PAD should be able to respond to a TEST frame received from a primary station DTE (including a PU type 2.1 that is acting as a primary station), as specified in the current IBM specification of SDLC.^[6]
6. RR, RNR, REJ, and I frames passed and Nr/Ns counts contained therein affect flow control and acknowledgment conditions on the local SDLC link. These acknowledgments and changes in flow control condition should be propagated to and reflected in the associated virtual circuit. When more than one virtual circuit is derived from a single SDLC link, only those frames with SDLC station addresses being mapped to a virtual circuit should have flow control or acknowledgment impact on that virtual circuit. See section 6.2 for additional details.
7. The information fields of I frames received from the DTE over the local SDLC link are mapped into the user data field of (the equivalent of) an X.25 unqualified Data packet (or packet sequence) on the virtual circuit corresponding to the local SDLC station address. Similarly, the user data field of incoming X.25 Data packets (or packet sequences) are mapped into the information field of an I frame with a corresponding SDLC station address. This mapping incorporates any needed segmentation/recombination and corresponding use of the M-bit. See sections 5.1 and 6.1 for details on SDLC/VC mapping procedures.
8. When error or failure conditions occur on the local SDLC interface that cannot be recovered locally at the link layer, treatment of the corresponding virtual circuit(s) should be handled in accordance with the procedures of section 6.4.

3.2.2.2 QLLC Support

For type A connections, the PAD supports SDLC over the interface to the local DTE and QLLC across the virtual connection to the remote end DTE or PAD. The SDLC PAD maps a subset of SDLC frames from and to corresponding QLLC frames carried in the user data field of qualified Data packets. Each qualified Data packet is carried over the virtual circuit whose logical channel number maps to the SDLC station address of the corresponding SDLC frame. Responses to those SDLC commands that are mapped to QLLC are delayed until corresponding responses from the remote end of the virtual circuit are received via QLLC Data packets. Those SDLC frames that are not mapped to or from QLLC qualified Data packets are handled at the local

SDLC interface, as is done for the Type B connection support option described earlier.

1. The SDLC PAD should support the QLLC configuration option. When this option is active, the following apply.
2. The PAD interface to the DTE conforms to section 3.2.1.
3. Handling of XID frames across the local SDLC interface, conforms to section 5.3, with the presumption of QLLC support.
4. The SDLC frames that are mapped to and from QLLC qualified Data packets are: SNRM(E), SABM(E), DISC, XID, TEST, UA, RD, DM, and FRMR. The RR frame can be mapped to and from QLLC QRR packets and the PAD should be able to receive and appropriately respond to such QRR packets. However, the PAD should not map RR frames received from the locally attached DTE to QRR packets to be sent over the virtual circuit connection to the remote end. Such mapping would be used to implement end-to-end polling across the virtual circuit, which is not used by the SDLC PAD because of the protocol overhead it entails (see section 3.2.2.3). The information fields of I frames are mapped into the user data field of unqualified Data packets in the same way as is done for the Type B connection option. All other SDLC frames are handled locally, as done for the Type B connection option. Section 8 of Part II of the IBM document specifying the QLLC protocol and IBM X.25 interfaces^[9] and Section 6.1.1 of this document provide additional details on the procedures associated with QLLC and the mapping between local SDLC frames and QLLC qualified Data packets on the corresponding VC.
5. When error or failure conditions occur on the local SDLC interface that cannot be recovered locally at the link layer, treatment of the corresponding virtual circuit(s) should be handled in accordance with the procedures of section 6.4.

3.2.2.3 PAD Polling Functions

For both the Type B connection and QLLC interface options, the SDLC PAD handles polling functions locally on the SDLC interface to the DTE. The PAD responds to polls from a directly connected primary station on behalf of the secondary station on the remote end of the virtual circuit connection. The PAD polls a directly connected secondary station on behalf of the primary station on the remote end of the virtual circuit connection. None of these local polls are propagated across the virtual circuit connection (e.g., in a QRR packet with the P bit set). The use of local polling avoids inefficient use of virtual circuit capacity for passing polls on an end-to-end basis.

Local port buffers (i.e., SDLC secondary station and virtual circuit send queues) on the PAD are used to store mapped/mappable frames and user information fields until they are passed to the appropriate local SDLC secondary station or virtual circuit connection. Normal X.25-based acknowledgement and flow control procedures are used over the virtual circuit connection to empty PAD port buffers. Local SDLC interface flow control, virtual circuit connection link and network layer flow control, and buffer size management are all coordinated by the PAD to achieve efficient transmission of data and avoid any deadlock conditions.

A poll consists of an SDLC command frame with the poll (P) bit set. A response consists of a sequence of frames from the secondary station addressed by the poll command, the last (and possibly only) of which has the final (F) bit set. A primary station may have only one outstanding poll on the interface that has yet to receive a valid response, within the specified time-out period.

1. The SDLC PAD should provide the following basic polling functions:
 - a. Assure that all active stations are polled at least once during each polling cycle.
 - b. Provide a mechanism to permit different stations on a line to be assigned different polling frequencies or priorities.
 - c. Provide a mechanism to permit inactive stations to be polled less frequently outside the period of modified operation defined in items 2 and 3 of section 6.3, and without an attempt to initially transmit I frames, to reduce unnecessary link overhead.

Following the numbered requirements list is the description of a specific polling approach which is designed to achieve these basic polling functions. The detailed procedures specified below illustrate the intent of the basic polling functions and represent a suggested implementation approach.

2. PAD buffers provided and the implementation of the local polling cycle should be such that the cycle is completed in no more than 0.25 second for the prototype cases specified in Appendix B, given that each secondary station has one polling slot in the cycle.

For each active secondary station on each SDLC interface, the PAD maintains a send queue. For an interface at the secondary station end, each send queue stores command frames to be sent across the local SDLC interface to the corresponding secondary station address. When the PAD SDLC interface is at the primary station end, each send queue is used to store response frames from the corresponding secondary station at the remote end of the virtual circuit, to be passed to the primary when the PAD is polled for that station. It is desirable that frames, added to a queue before the (previously) last frame is sent, update the "last frame" used in determining when the P/F bit is set. These SDLC interface queues are separate from the logical channel send queues also maintained by the PAD for storing packets to be sent into the network over the corresponding virtual circuit. Section 6.2 addresses procedures used by the PAD to maintain send queues.

For SDLC interfaces to primary stations, the primary station is responsible for local polling and the PAD interface responds on behalf of actual or virtual secondary stations for which the line is configured, in accordance with SDLC procedures described and referenced above. When a secondary station is polled, the PAD responds with the contents of the corresponding send queue or with an RR or RNR if the queue is empty. The F bit of the last (or only) frame is set.

For SDLC interfaces to secondary stations, the PAD takes the active role in polling each secondary station. In this case, the PAD will execute a poll loop covering all secondary stations on the line. The type of polling treatment an individual station receives depends upon its current polling status.

At any time an operational station will be considered either to be "active" or "inactive" by the PAD. All operational stations associated with a virtual circuit are considered to be "active" when the SDLC interface is initialized, following a valid UA response to a SNRM(E)/SABM(E). An active station becomes inactive after the PAD fails to receive a valid response to a poll within the number of retries configured on the interface for this purpose. An inactive station becomes active following a valid response to a poll from the PAD within the time-out period. A station is not operational for polling purposes if it has been declared "nonoperational" or is not associated with an established VC (for a reason other than local end station inactivity).

The PAD should clear an established SVC or reset the PVC (see item 2 of section 5.4) if the corresponding secondary station becomes inactive. In addition, the PAD should not permit an

SVC to be established to or from a secondary station that is inactive or is not in the NRM or ABM mode. The PAD should establish an SVC if the interface is configured for auto call service and the secondary station returns to active status.

The polling of stations during a cycle is governed by a set of configurable parameters:

- The number of polling slots per cycle (Nps)
- The polling vector for each cycle (list of stations, with repetitions used to achieve differential station priorities, for the Nps polling slots)
- The number of slots to be skipped for an inactive station before it is actually polled (activation response-lag factor, ARF)

For each polling slot of a cycle, the PAD polls the corresponding station if the station is operational and a corresponding VC is established. However, if the station is inactive and the period of modified operation defined in section 6.3 has expired, the PAD decrements the ARF counter for each polling slot until the zero value is reached (at which point the counter is reset and the station is polled). Polling is in all cases subject to flow control constraints, as specified below. For active stations that are not flow controlled, the poll consists of a sequence of I frames queued for the station, the last of which has the poll bit set, or an RR frame with the poll bit set if no frames are available for transmission to that station. For inactive stations, the poll consists of a SNRM(E) or SABM(E) with the poll bit set.

The PAD then awaits a valid frame sequence response. (While waiting, the PAD may transmit I frames on a non-poll basis to other secondary stations on the interface.) Following a valid response sequence, the next slot in the cycle is polled. If a valid response is received from an inactive station, it is restored to active status, and the PAD polls it with an I-frame sequence or RR frame, as appropriate. If a valid response is not received from a currently active station after the specified number of retries, it is placed into the inactive status.

If the PAD is flow controlling the station corresponding to a polling slot, it either skips that slot if there is no I frames to send, or it sends a sequence of frames to the station ending with an RNR frame with the poll bit set. If the PAD is being flow controlled by the station corresponding to a polling slot, it sends an RR or RNR frame, as appropriate, with the poll bit set, and awaits a valid response. A flow control condition exists if the end of the unacknowledged transmission window has been reached, or an RNR frame has been received from the remote end of the SDLC link and this busy condition has not yet been cleared.

The above described polling procedure is illustrated in Figure 3-1 with an SDL (System Description Language) flowchart. This procedure represents one possible approach, intended to achieve the basic polling functions cited in requirement item 1 above.

Table 3-1. Summary of Physical and Link Layers of PAD SDLC Interface Support

Physical Level

Line Speeds (bps)
Non-switched: 300, 1200, 2400, 4800, 9600, 19200, 56000
Switched (dial): 300, 1200, 2400, 4800, 9600

Interface Standard
EIA 232D: Line speeds of 19200 bps or less
CCITT V.35: Line speeds greater than 19200 bps
X.21*

Duplex and half-duplex operation options

Link Level

Normal Response Mode (NRM) & Combined (NRM/ABM-capable) Options
Modulo 8 Frame Sequencing
Modulo 128 Frame Sequencing*
Link Level Window Sizes: 1-7 (mod. 8); 1-60, 61-127* (mod. 128*)

Switched Interfaces
Point-to-point operation

Non-Switched Interfaces
Point-to-point and multipoint operation options

Two-Way Alternate & Two-Way Simultaneous Operation Options

Minimum Number of Stations Supported Per SDLC Interface
2400 bps or less: 4
4800-9600 bps: 16
greater than 9600 bps: 32

Support of primary, secondary, or negotiable local interface station options

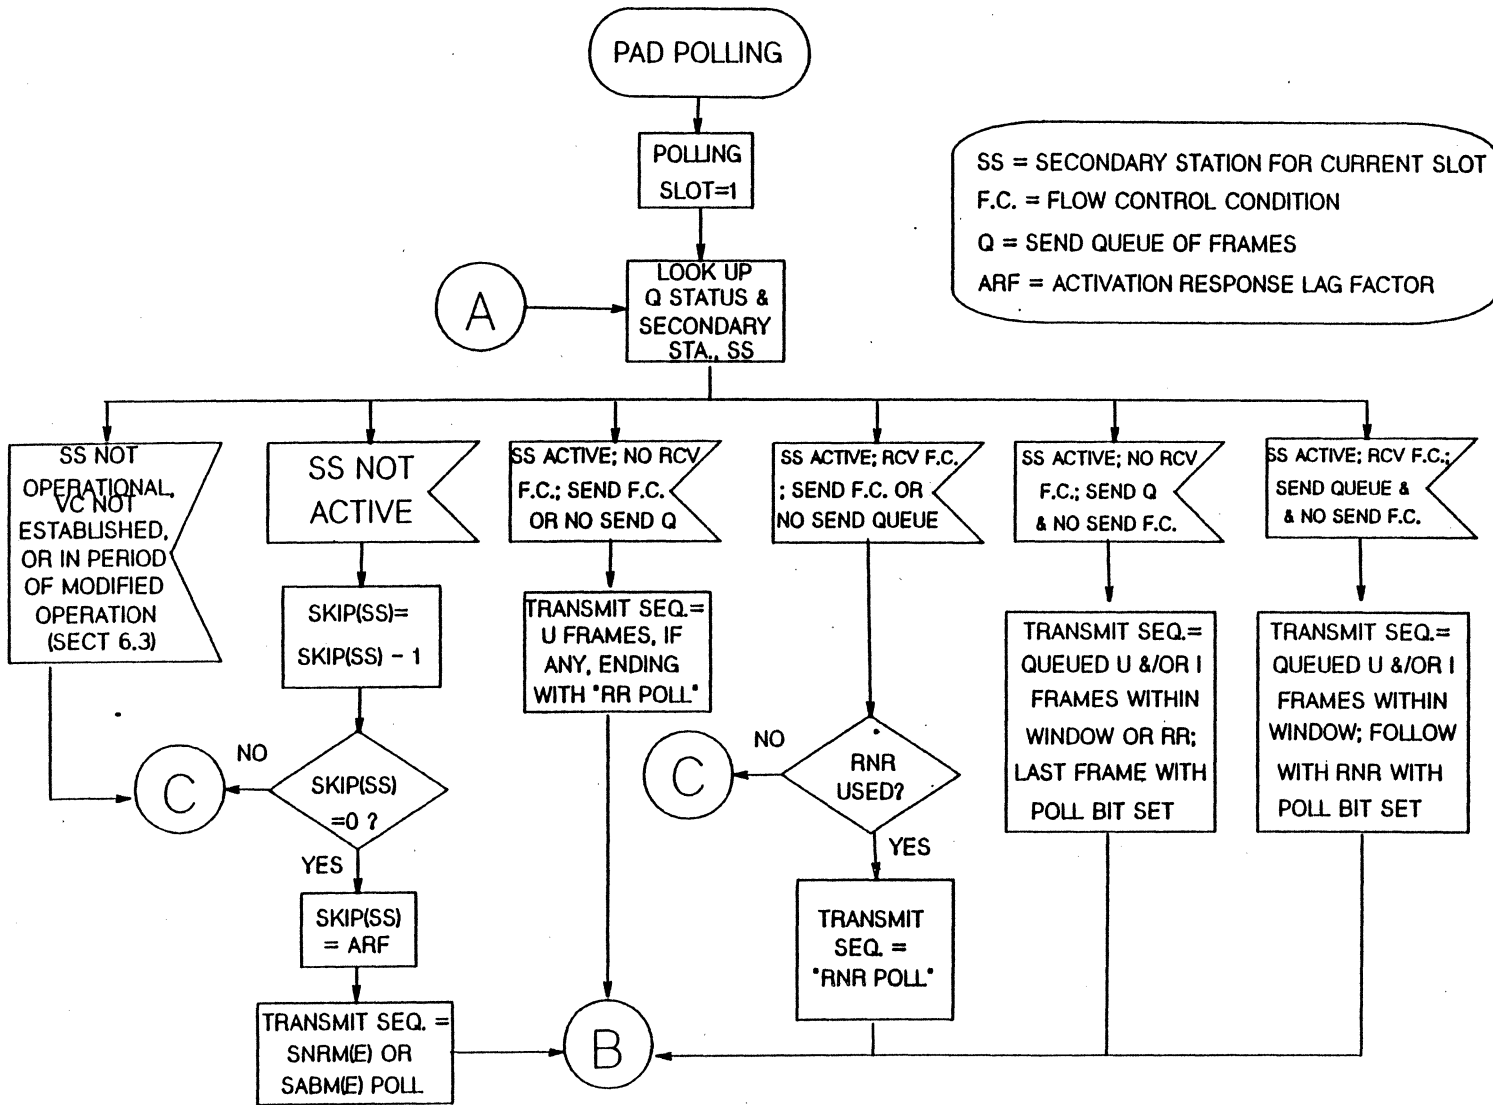
Max. Frame Size to Accommodate I Frame Information Field of:
128 to 137 octets
256 to 265 octets
512 to 521 octets*

Dynamically Settable By XID Exchange on Switched Interfaces:
ABM capable or not
2-way alternate vs. 2-way simultaneous operation
primary, secondary, or negotiable local station
maximum I frame information field length
link level window for each secondary station

All SDLC Frame Types Except
RIM, SIM
UP, CFGR, BCN
UI

Note: All items listed above should be supported, unless starred (*).
Support of starred items is desirable.

Figure 3-1. SDLC Flowchart of SDLC PAD Polling of Secondary Stations



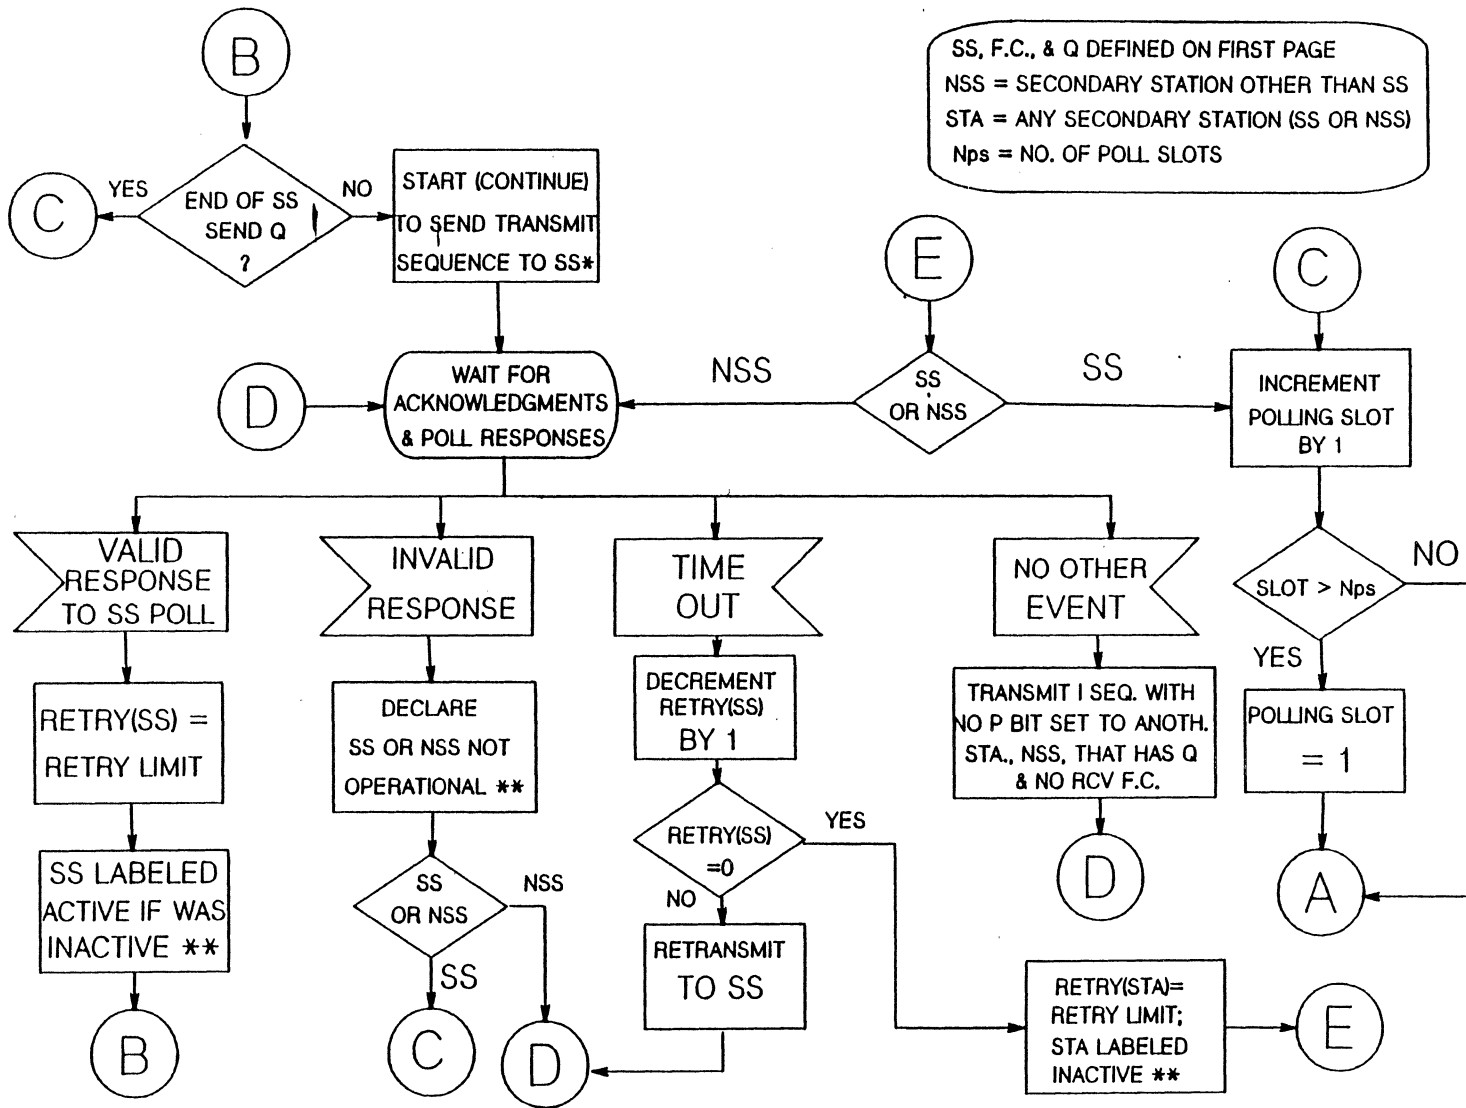


Figure 3-1. SDLC Flowchart of SDLC PAD Polling of Secondary Stations (Cont.)

* UPDATE SEQUENCE DYNAMICALLY TO INCLUDE ANY RECENTLY RECEIVED FRAMES FOR SS.

** ALSO TAKE ANY APPROPRIATE VC ESTABLISHMENT, CLEARING, OR RESETTNG ACTION IN RESPONSE TO ANY STATUS CHANGE.

BLANK PAGE

4. PAD Virtual Circuit Support

4.1 Support of PVCs and SVCs

For each active SDLC interface that a PAD terminates locally, one or more virtual circuits are maintained within PPSN toward the remote end DTE(s). If the PPSN network element containing the SDLC PAD connects to a network element within the BOC network of another manufacturer in establishing the virtual circuit, a logical channel on an X.75' interface between the two network elements is used. Otherwise, the SDLC PAD may maintain the virtual circuit by means of proprietary protocol interfaces with other network elements. If the remote end SDLC PAD, X.3/X.28/X.29 PAD, or X.25 DTE is not on the same network, the respective virtual circuit will "leave" the network across an X.75 interface.

1. In all cases, the virtual circuit supported locally by an SDLC PAD should be compatible with any intervening X.75 or X.75' interfaces in the call path and should permit successful communication (up to and including the packet layer) with a 1984 X.25 DTE or a DTE served by an X.3/X.28/X.29 or PPSN SDLC PAD at the remote end. Any X.75' interfaces on a network element containing an SDLC PAD should conform to sections 3.4 and 3.5 of the PPSNGR.^[3]
2. An SDLC PAD should be able to simultaneously support at least a minimum number of logical channels for each configured local SDLC interface. This minimum number of logical channels per SDLC interface varies with the line speed of the SDLC interface, as follows:

SDLC Line Speed	Min. # of LCs per SDLC
2400 bps or less	4
4800-9600 bps	16
over 9600 bps	32

Also see section 5.1.1 for requirements on the relationship between the maximum number of secondary stations that can be configured and the number of simultaneous virtual circuits that should be supportable.

3. Within the overall network element or SDLC PAD resource constraint on the total number of logical channels that can be supported, it should be possible to configure the SDLC PAD to support any number of switched virtual circuits (SVCs) and any number of permanent virtual circuits (PVCs).
4. When the connection is an SVC, the SDLC PAD originating a Call Request should be capable of populating the CUD field of the packet as specified in item 3 of section 5.4.2.1. When receiving a call request at the called end, the SDLC PAD should be capable of correctly interpreting the contents of the CUD field, as specified in item 3 of section 5.4.2.1. If an invalid CUD value is received in an incoming call request, the PAD should clear the call with cause "DTE originated" and diagnostic code decimal 235.
5. Support of SVC and PVC service by the PAD for Type A connections should conform to the IBM QLLC specifications.^[9] In particular, the PAD should:
 - support the Q bit mechanism for identifying QLLC commands and responses in Type A connections

- support the M bit mechanism
 - support both values of the D bit on Type B connections but use only the "0" value of this bit on Type A connections
 - always include the diagnostic code field in Clear Request, Reset Request, and Restart Request packets
 - support the diagnostic codes specified in Appendix F of the above cited IBM document when the cause is "DTE originated" [when the SDLC PAD generates the clear, reset, or restart request acting as an X.25 DTE on behalf of the actual SDLC DTE, it will use the DTE originated cause and diagnostic codes; for example, when an invalid protocol identifier is detected in the CUD – cause code decimal 0 and diagnostic code decimal 235]
 - never generate an Interrupt packet, but respond with an Interrupt Confirmation if an Interrupt packet is received (the interrupt user data being ignored)
 - conform to section 4.4.3 in part II of the IBM specifications for SNA/X.25 DTE/DCE interfaces^[9] in the generation and handling of SVC and PVC resets
6. The features, options, and parameter values that should be supported by SDLC PADs for SVC and PVC connections are specified in Table 4-1. Unless otherwise specified, these features, options, and parameters are as defined in the 1984 version of CCITT Recommendations X.25 and X.75 and in sections 3.2 and 3.4 of the PPSNGR.
 7. Each SDLC PAD should support two additional subscription options for each configured SDLC interface using SVCs: (1) Call Redirection Acceptance and (2) Called Line Address Modified Acceptance. The first controls whether or not the PAD should clear call requests in which the local interface is an alternate destination, following call redirection. The second controls whether or not the PAD should clear call connection indications in which the called address is not the same as the originally called address of the call request. Since the actual user will probably not have access to the redirection or address change notification, these subscription options permit users to prevent such changes from occurring. The end of the virtual circuit remote from the clearing PAD should receive the clearing cause "DTE originated" and diagnostic code #240 (remote procedure error). In the case of clearing because of non-acceptance of Called Line Address Modified Notification, the end of the virtual circuit local to the clearing PAD should receive the same clearing cause and diagnostic code.

4.2 X.25 Product Compatibility

The PPSN supports DTE's conforming to the 1984 version of CCITT Recommendation X.25 as set forth in PPSNGR Section 3.2 (or the 1980 version, as a configuration option). Thus, PPSN X.25 DTE/DCE interfaces can be expected to be compatible with X.25 products that conform to these versions of CCITT Recommendation X.25. The following illustrates some of the SNA products providing such interfaces:

Products Supporting NPSI
3705 (Supports PVCs only)
3720, 3725, and 3745 (Supports PVCs and SVCs)

Integrated X.25 Products
Series/1
System/36
System/38
3174
3274
3710
5973
8100 DPPX (not DPCX)
9370

Note: The above illustrations are model designations for IBM products, but other manufacturers provide products with X.25 interfaces that can be used in an SNA environment.

However, efficiently servicing SNA equipment over a PSPDN packet switched connection requires that careful attention be paid to several tuning considerations. These considerations affect the performance of NPSI with X.25 networks. Tuning applies to several SNA buffer sizes and X.25 parameter values, each assigned for operation in their SNA and X.25 network environments, respectively. The SNA values, when correctly set, decrease the number of packets being generated while transmitting the same amount of user data. The proper X.25 parameters decrease the need for acknowledgements and, as a result, improve throughput.

Thus, proper configuration of NPSI (or other SNA X.25 interface products) and each connecting X.25 network DTE/DCE interface can improve performance (for example, see IBMs "Tuning Considerations for IBM SNA X.25 DTE's^[10]). Configuration of end user CPE and DTE/DCE interfaces is not generally a matter of network equipment generic requirements. However, the provision of special X.25 interface default configurations appropriate for SNA X.25 DTEs may be a useful network feature. The need for such SNA DTE default configuration options is under consideration.

Table 4-1. SDLC PAD Support of SVCs and PVCs

Packet Level (For SVCs and PVCs)

Modulo 8 Packet Sequence Numbering
Modulo 128 Packet Sequence Numbering*
Min. Logical Channels Per Configured SDLC Line, by Line Speed
 2400 bps or less: 4
 4800-9600 bps: 16
 over 9600 bps: 32
Window Size: 1 to 7 (modulo 8); 1 to 60 (modulo 128)*
Throughput Classes, by Line Speed
 9600 bps or less: all defined values up to line speed
 over 9600 bps: 9600 bps, 19200 bps*, 48000 bps*
Maximum Packet Sizes: 128, 256, 512*
Octet-Aligned User Data Field
Support of M Bit
Support of D Bit (Value of 1 for Type B Connections Only)
Support of Interrupt Packet (For Type B Connections Only)

CCITT Subscription Facilities Applying to SVCs to/from an SDLC Interface

Incoming Calls Barred
Outgoing Calls Barred
One-Way Logical Channel Outgoing
One-Way Logical Channel Incoming
Nonstandard Default Window Sizes
Nonstandard Default Packet Sizes
Flow Control Parameter Negotiation
Default Throughput Classes Assignment
Throughput Class Negotiation
Closed User Groups:
 CUG, CUG/OA, CUG/LA, CUG/OA/LA
 Up to 9600 bps SDLC interfaces: 10 CUGs/interface
 Over 9600 bps SDLC interfaces: 100 CUGs/interface
Fast Select Acceptance* (For Type B Connections Only)
Reverse Charging Acceptance
Local Charging Prevention
Network User Identification Subscription
Network User Identification Override
Hunt Group
 10 interfaces
 single address type
Call Redirection

Note: All items listed above should be supported, unless starred (*).
Support of starred items is desirable.

Table 4-1. SDLC PAD Support of SVCs and PVCs (Cont.)

PPSN Subscription Facilities Applying to SVCs to/from an SDLC Interface

- IC Preselection
- RPOA Selection Barred
- Automatic Call Service (see sect. 5.4.2.2)
- Virtual Circuit Profiles (see sect. 5.5)
- Call Redirection Acceptance
- Called Line Address Modified Acceptance

CCITT Per-Call Facilities

- Closed User Group Selection
 - Basic Format CUG
 - Extended Format CUG (> 9600 bps)
 - CUG/OA Selection
- Fast Select* (For Type B Connections Only)
- Reverse Charging
- Network User Identification Selection
- RPOA Selection
 - Basic Format
 - Extended Format*
- Flow Control Parameter Negotiation
- Throughput Class Negotiation
- Call Redirection Notification
- Called Line Address Modified Notification
- Transit Delay Selection and Indication

Note: All items listed above should be supported, unless starred (*).
Support of starred items is desirable.

BLANK PAGE

5. Associating SNA/SDLC PUs with Virtual Circuits

Given the virtual circuit service provided by SDLC PADs, as described above, this section addresses how sessions involving SNA PUs and LUs served off a PAD-terminated SDLC line are mapped into and carried over PPSN virtual circuits, at the service and concept level. Section 6 focuses on details of state transitions and frame and packet mapping for the SDLC and X.25 packet layer protocols involved.

5.1 Mapping LUs & PUs to and from VCs

5.1.1 General Requirements

When connecting a local SDLC interface to the remote end PAD or DTE via a virtual circuit connection, it is necessary to associate each SDLC station with a separate logical channel number. Information associated with different SDLC stations (SNA PUs) must be carried within the network on separate virtual circuits for two reasons:

- The station address distinction available in the SDLC protocol is not available in the link level service (i.e., LAPB) provided over X.25 virtual circuits. The LAPB protocol permits only point-to-point link configurations. Thus, for example, an SDLC I-frame mapped into a LAPB I-frame would no longer contain an address capable of distinguishing among different stations for a given logical channel number.
- It is important to be able to simultaneously connect different stations on an SDLC link with different remote end DTEs. This can be done through the packet network only if each station is associated with a separate virtual circuit.

Thus, each virtual circuit (VC), identified locally via a logical channel number (LCN), should be associated with a connection between one SDLC station at one end and either one SDLC station or an X.25 virtual circuit termination at the other end. The only question is whether the SDLC PAD should be able to support more than one virtual circuit for a given local SDLC station. Specifically, should the PAD be able to associate groups of one or more LUs on a station (PU) with their own virtual circuits.

This SDLC capability would permit different display terminals or subgroups of such terminals on a single cluster controller to simultaneously connect to geographically separated hosts via separate VCs, rather than relying upon SNA routing capabilities associated with LU-LU session establishment. Since this capability would require PAD support of layers of the SNA protocol stack above SDLC, it could also open the possibility of PAD support of terminal re-configuration without the need for ACF/NCP and ACF/VTAM changes. In the interface with an SNA front end processor (FEP), the PAD could maintain static virtual terminal (LU) or terminal pool configurations. Reconfigurations at the terminal end could then be handled via PPSN translations rather than FEP and host reconfigurations.

However, this LU mapping capability requires support of higher layers of the SNA protocol structure. The PAD would have to act as an SSCP and understand the SSCP-to-PU and SSCP-to-LU session protocols to achieve this functionality. In addition, maintenance of customer network management functions in an environment with LUs of a single station being assigned to different virtual circuits may require significant additional PAD functionality. As mentioned above, IBM SNA itself provides the ability for different LUs within the same PU to simultaneously connect to geographically separated host applications as part of its routing capabilities for LU-LU session establishment. Figure 5-1 illustrates how SNA would permit three terminal devices on a single cluster controller in one region of the country to simultaneously connect (with a VC replacing a dedicated long-haul SDLC link) to applications in three

geographically separated hosts in another region of the country. Depending on the traffic levels between the host on the end of the long-haul virtual circuit and the ultimate destination host, the final leg of the connection could be via a direct SDLC link between FEPs or a local VC replacement for such a link (cases B and C of the Figure, respectively).

Because of the significant increase in protocol and buffer complexity required for mapping LUs to VCs in the PAD and the availability of routing capabilities within SNA to permit LUs to connect to geographically separate applications, whether or not this additional PAD capability should be supported is under consideration.

1. The SDLC.PAD should support the ability to associate (on a one-to-one basis) each actual or emulated secondary station (SNA PU) on each local SDLC interface with a virtual circuit connection. This virtual circuit (PVC or SVC) is identified at the PAD interface by its local logical channel number (LCN). Packets sent across the virtual circuit, that are mapped from frames received by the PAD over the local SDLC link, should contain the LCN currently associated in PAD memory with the station address of the SDLC frame. Similarly, frames passed to the DTE over the local SDLC interface as a result of a packet received on a virtual circuit should contain the station address currently associated in PAD memory with the LCN of that packet.
2. The SDLC concept of group address need not be supported by the SDLC PAD at this time. Possible PAD support of SDLC group addresses through the use of address mapping and mapping of frames to multiple associated virtual circuits is under consideration.
3. The SDLC concept of broadcast address should be supported by the PAD. In the point-to-point configuration, a broadcast address is equivalent to the address of the only permitted secondary station.
4. The SDLC PAD should be able to simultaneously support at least N_{vc} virtual circuit connections on the PPSN associated with SDLC interfaces, in addition to virtual circuits needed to support non-SDLC access interfaces on the PAD/AC. N_{vc} is the maximum number of (actual or emulated) SDLC secondary stations that can be configured for the PAD.
5. Each PAD SDLC interface should include among its configuration parameters, one that specifies if the interface is switched (dial-up), non-switched with "fixed" (i.e., actual) secondary stations, or non-switched with "grouped" (i.e., virtual) secondary stations, as described below.
6. Table 5-1 identifies which combinations of SDLC PAD interface configuration options are valid and which are not, given current requirements.
7. In mapping between SDLC frames and VC packets (especially when SDLC responses to local SDLC commands must await QLLC packet responses from the remote end of the virtual circuit) SDLC procedures concerning the order of frame responses on the local interface should not be violated by the PAD. If necessary, the PAD should discard QLLC responses under timeout conditions to avoid protocol violations or frame response ambiguity/redundancy on the local SDLC interface.

Additional requirements pertaining to the mapping of PUs to and from virtual circuit connections depend on the nature of the SDLC interface the PAD port is supporting. Three categories of SDLC interface PAD ports are distinguished for this purpose:

- Ports for non-switched SDLC interfaces dedicated to PUs and LUs that are fixed in number and type by interface configuration parameters. [Please note that as used in this context,

"dedicated" means reserved, rather than non-switched.] Such ports serve non-switched interfaces to secondary stations, and to PU type 2.1 stations and primary stations (e.g., host access ports) that are dedicated to packet network access by a fixed, remote secondary station configuration or users whose PU/LU configurations do not differ appreciably.

- Ports supporting switched (dial-up) SDLC interfaces.
- Ports for non-switched SDLC interfaces serving virtual PUs and LUs whose numbers and characteristics are designed to best match the mix of actual PUs and their LUs that are expected to access this port from the remote end of the virtual circuit. These ports can be considered to have "grouped" PU/LU characteristics. Such ports serve non-switched interfaces to primary stations that expect to be accessed by a diversity of users with varied PU/LU configurations.

The next three subsections present additional requirements appropriate for each of these three PAD SDLC support situations, respectively.

5.1.2 Requirements for Ports with Fixed PU/LU Characteristics

The simplest case for a non-switched SDLC port is that which is configured to serve a known PU/LU configuration of secondary stations or a primary station which serves as a host access port dedicated to a particular remote user with a fixed PU/LU configuration of secondary stations. At the terminal end, such a PAD interface would normally be appropriate, since the configuration of control units and attached devices would be known and fairly stable for non-switched access. At the host end, such a PAD interface would be appropriate for an SDLC access line dedicated to a particular remote user or a limited selection of users with identical or very similar PU/LU configurations.

For this case, the SDLC interface assumes a specific number of PUs. Each PU is assigned a specific number of LUs. Those other PU and LU characteristics which must be specified as NCP and/or VTAM parameters in configuring the host end access port on the FEP are also fixed for this case. When used at the host end, the PAD is emulating actual secondary stations at the remote end of the virtual circuit. Although this case is not appropriate for host access by varied users, it provides a close fit between the PUs and LUs emulated by the PAD at the host end and the actual PUs and LUs making up the secondary stations at the terminal end of the virtual circuit connection. Thus, there is little or no waste of PAD or host FEP resources in configuring and emulating PUs or LUs that are not really present at the terminal end.

Because each PU is mapped to a separate virtual circuit and the multi-point station address mechanism of SDLC is not available within the X.25 link layer service, it is necessary to assign separate X.121 DTE addresses to each PU on a fixed characteristics interface. At a PAD servicing secondary stations across the SDLC link, the SDLC station addresses are those of the actual PUs. At the PAD servicing the primary station across the SDLC link, the SDLC station addresses are those of the PUs being emulated within the PAD on behalf of the remote end secondary stations.

1. For SDLC PAD interfaces configured as non-switched and serving virtual circuit connections with "actual" secondary stations at the terminal end, the number of PUs (secondary stations) and their station addresses are configuration options. These would be configured to match the values included in the NCP/VTAM specifications used to construct the SDLC access interface at the host end.
2. For such interfaces, the PAD will support separate X.121 addresses for each PU. These addresses pertain to the secondary stations actually connected to the PAD or being emulated by the PAD. The PUs may also be members of hunt groups, if hunt

groups are supported for fixed characteristics ports. See section 5.2 for additional details.

3. Non-switched, fixed characteristics SDLC interfaces should be capable of serving as either end of a PVC or SVC connection.

The number of LUs associated with each PU and LU/PU characteristics other than those already specified as PAD interface configuration options are reflected only in the NCP and/or access method specifications at the host end. For example, the number of LUs per PU and whether a PU is type 1 or type 2 is of no concern to the PAD, unless the (under consideration) ability to map LUs to separate virtual circuits is implemented.

5.1.3 Requirements for Dial-Up (Switched) Ports

For switched SDLC interfaces, only the point-to-point configuration is supported. Thus, there can be only one secondary station. In this case, the number of virtual circuits associated with an active interface is the same as the number of secondary stations: one.

For switched SDLC interfaces, NCP/VTAM permit the PU type to be specified as "either 1 or 2" (the "either" designation is one of the options), which would be the safest designation at the host end unless it were certain that only secondary stations of a single PU type were to access the host port. In addition, for switched access at the host end, LUs associated with the secondary PU are treated on a pooled, rather than explicitly designated, basis. LU characteristics are ascertained dynamically after the switched connection is established (by the associated SNA session establishment activities). The disadvantage of a switched connection is that only one secondary station can be supported on the interface. Thus, a PAD cannot concentrate traffic from multiple secondary PUs on switched SDLC ports on a FEP at the host end.

At this time, the greatest need for switched SDLC interfaces is considered to be dial-out by the PAD at the secondary station end. PPSN support of other SDLC switched interface options is under consideration.

Since switched SDLC interfaces may only be point-to-point, it is only necessary to assign one X.121 address per SDLC interface.

1. For SDLC PAD interfaces configured as switched, only one secondary station is permitted, and its station address is a configuration option.
2. For switched SDLC interfaces, the PAD will assign an X.121 DTE address to that interface. See section 5.2 for additional details.
3. The PAD should be capable of supporting dial-out local SDLC interfaces at the secondary station end. This switched SDLC interface should be capable of serving as either the calling or called end for SVC service. Support of SDLC interface dial-in and switched SDLC service at the primary station end are under consideration.

5.1.4 Requirements for Ports with Grouped PU/LU Characteristics

It may not always be possible or economical to dedicate non-switched SDLC access ports at the host end to a single remote user at the terminal end. For example, it may be desirable to offer remote access to some host applications from a selection of geographically dispersed terminal equipment with varying PU/LU characteristics. If the duration or amount of access by such terminal equipment sites does not justify establishing a non-switched port dedicated to specific terminal equipment at the FEP, a more flexible alternative would be appropriate. To accommodate such situations within the constraints of IBM SNA access port configuration rules (i.e., NCP and access method), the concept of a non-switched interface with "grouped PU/LU characteristics" is introduced for use on the host (primary station) end of the virtual circuit

connection.

As with the fixed PU/LU characteristics option, the SDLC interface assumes a specific number of PUs and each PU is assigned a specific number of LUs. This is required by the way in which IBM FEP access ports are configured. Those other PU and LU characteristics which must be specified as NCP and/or VTAM parameters in configuring the host end access port on the FEP are also fixed, as for the "fixed" case. However, the way in which these configured PUs and LUs are used is different in the "grouped" case.

When a grouped PU/LU characteristics access port is used, the PAD connected to the host end is emulating virtual secondary stations – secondary stations that are representative of those expected to actually access the host port from the remote terminal end. To accommodate the diversity in PU/LU characteristics of the terminal equipment actually expected to access the port, a grouping strategy is adopted, which differentiates this type of port from the "fixed" type. Each PU configured on the interface is intended to be representative of a class of potential terminal end users rather than a specific remote end user. The number of PUs configured, the number of LUs configured for each PU, and the PU/LU characteristics are designed to accommodate the expected access traffic load for the host end interface, without creating a large number of virtual PUs and LUs at the host end which rarely have an actual counterpart at the terminal end. The PUs configured at the host end are assigned characteristics that are representative of a group of similar terminals expected to access the interface.

Three factors enhance the viability of this grouped approach:

- The number of PU and LU characteristics that must be specified as part of the FEP access port configuration is limited. These would be the characteristics used for "grouping" the virtual station PUs and LUs at the host end interface. Specifically, the "groups" for such an interface would be defined as follows. Each PU is categorized as either a type 1 or type 2. For each PU or group of PUs within the above two categories, a subgroup would be defined in terms of the number of LUs that use formatted commands and the number of LUs that use character-coded commands when communicating with the SSCP (VTAM LU operand SSCPFM). Other PU and LU characteristics are determined in a fashion that is transparent to the PAD interface (e.g., the BIND for a LU-LU session).
- Although the number of PUs and LUs is fixed for non-switched interfaces on a FEP, not all of them must have an active counterpart at the terminal end. The PAD must emulate each of these at the host end, but this emulation hides from the host the fact that not all PUs and LUs configured have an actual active control unit or device at the remote end. The key is to have enough PUs and LUs per PU to meet demand most of the time, without expensive over-engineering of this capacity.
- The dynamic reconfiguration capability of NCP/VTAM is used to reduce the difficulty of making configuration changes within the customer's IBM host/FEP equipment to accommodate changes in the number or type of terminal equipment expected to access the port. This further reduces the need to over-engineer the initial configuration.

When a virtual circuit connection is established from a secondary station to a grouped characteristics SDLC interface on a host FEP, the PAD at the host end chooses an available secondary station PU from among the configured virtual stations, based on the type of calling secondary station and the PU groups available. Within the PUs available on the SDLC interface(s) of the DTE or DTE hunt group of the specified host X.121 address (see section 5.2), the PAD selects an available PU that matches the type of the calling PU and has the smallest number of LUs of each type that is greater or equal to that specified for the calling secondary station. The information on the secondary station type is passed to the host end PAD in the call user data field,

following the protocol identifier (see sections 4.2 and 5.4.2.1).

Although grouped characteristics SDLC interfaces may be multi-point, it is not necessary to support separate X.121 addresses for each secondary PU associated with the interface. Whereas for the fixed characteristics SDLC interface, each secondary PU at the terminal end is mapped to a specific PU being emulated within a PAD at the host end of the virtual circuit connection, for this type of interface the PAD at the host end selects among the available PUs on the designated SDLC interface or hunt group as described above. Since the PAD uses a logical procedure to select the correct host end secondary PU emulation (i.e., SDLC station address), it is only necessary to route the virtual circuit connection to the correct SDLC interface. Thus, on grouped characteristics interfaces, X.121 DTE addresses are assigned only on a per-interface basis. Such interfaces may be members of hunt groups.

1. For SDLC PAD interfaces configured as non-switched and serving virtual circuit connections with secondary stations at the terminal end that may vary over time, the number of PUs (secondary stations) and their station addresses should be configuration options. These options are configured to match the values included in the NCP/VTAM specifications used to construct the SDLC access interface at the host end.
2. For these same SDLC PAD interfaces, each configured PU should be associated with a "group" or PU/LU characteristics class. Each group should be defined by the following in the configuration data:
 - whether the PU is type 1 or type 2
 - the number of LUs associated with the PU that use formatted commands in communicating with SSCPs
 - the number of LUs associated with the PU that use character-coded commands in communicating with SSCPs

The PAD should be able to support at least 16 different groups, definable by the BOC in terms of the above three indices. Support of a larger number of groups is desirable. The number of each of the two types of LU, per PU, should be settable to any value in the range 0 through 32.

3. If the SDLC interface(s) corresponding to the host end DTE address or hunt group is of the "grouped" characteristics type, then for a terminal end initiated virtual call, the host end PAD should assign the virtual circuit to the PU (SDLC interface and station address) as follows. Only a PU of the same type as the calling secondary station should be selected. Within this constraint, the virtual circuit is assigned to a PU with the smallest number of LUs which nevertheless satisfies the number of LUs of the two types specified for the calling secondary station. The allocation is made across all PUs on the specified SDLC interface or hunt group of such interfaces. In making this assignment, the group identifier of each PU and group definition in the PAD is used to determine host end SDLC characteristics and the call user data field of the Call Request packet provides the characteristics of the calling secondary station end. As each virtual circuit is assigned, the requested throughput class value is subtracted from the throughput value total configured for the corresponding SDLC interface. If the result is less than zero, the allocation is not made and the call is cleared as specified below. When a virtual circuit is cleared, the corresponding throughput class value is restored to the "remaining capacity" value for the corresponding SDLC interface. The throughput value total configured for the SDLC interface should be selectable from among at least the following fractions/multiples of the line speed configured

for the interface: 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0, 1.25, 1.5, 1.75, 2, 2.5, 3, 5, 10, and "no limit."

If no PU is configured which satisfies the above requirements, the call is cleared with cause "Incompatible destination" and diagnostic code #0. If at least one PU is configured which satisfies the above requirement, but none is currently available on an interface with adequate remaining throughput class capacity, the call is cleared with cause "Number busy" and diagnostic code #71.

4. The PAD supports separate X.121 DTE addresses for each grouped characteristics SDLC interface. Such interfaces may be members of hunt groups. See section 5.2 for additional details.
5. Non-switched, grouped characteristics SDLC interfaces should be capable of serving as either end (i.e., calling or called) of an SVC connection.

5.2 X.121 Address Assignments & Hunt Groups

The specificity of an X.121 DTE address depends on the type of SDLC port on which it is configured. An X.121 address is assigned to each switched and non-switched, grouped characteristics SDLC interface configured on a PAD. On non-switched, fixed characteristics SDLC interfaces, a separate X.121 address is assigned to each secondary station address (i.e., actual or emulated PU). These X.121 addresses are in accordance with the rules for PPSN addresses specified in section 4 of the PPSNGR. [If PAD support of multiple VCs per SDLC station (PU) is specified in a future issue of this document, the need to distinguish individual LUs or subsets of LUs within a PU in establishing a VC would have to be considered.]

Support of hunt groups also varies with the type of SDLC interface. Because of the one-to-one association of secondary stations with PU emulations supported on fixed characteristics SDLC interfaces at the remote (host) end of the virtual circuit connection, support of hunt groups is not vital for non-switched, fixed characteristics SDLC interfaces. Because switched SDLC interfaces are currently specified only for PAD dial-out to the DTE at the terminal end, hunt groups for switched interfaces are not considered necessary.

Non-switched, grouped characteristics interfaces require a special type of hunting, in which individual PUs are assigned from among all the PUs on all of the interfaces within the hunt group, based on the rules specified earlier. Because of the need to match actual secondary PUs at the terminal end with compatible virtual PUs being emulated within the host-end PAD, normal interface hunting algorithms are not appropriate.

1. Each SDLC interface on an SDLC PAD configured as switched or non-switched, grouped characteristics should be configured with a single X.121 DTE address.
2. Each secondary station configured on a non-switched, fixed characteristics SDLC interface should also be configured with a single X.121 DTE address. [Both the need for support of multiple VCs per secondary station and the approach for distinguishing LUs (separate X.121 DTE addresses or another mechanism) associated with such support are under consideration.]
3. The PAD should provide, as a subscription option, a hunt group capability that assigns incoming virtual calls containing a hunt group called address to a PU from among all PUs configured on all the non-switched, grouped characteristics SDLC interfaces that are members of the hunt group. The assignment is as described in section 5.1.4.

4. A collection of non-switched, grouped characteristics SDLC interfaces may be configured as a hunt group. Such a hunt group should be able to be configured with at least 10 member SDLC interfaces within the same AC or PS.
5. Only the single address hunt group type, as defined in section 4 of the PPSNGR, should be supported.
6. It should be possible to add, activate, deactivate, and delete interfaces in a hunt group without affecting services on other interfaces in the hunt group.
7. If an interface within the hunt group becomes nonoperational, any virtual call associated with the interface should be cleared. The interface should be declared unavailable and no new call setups should be directed to that particular interface.
8. If a hunt group has individual addressing (i.e., an individual SDLC interface has an individual address distinct from the common address of the hunt group), incoming calls to the individual address should be hunted only within the PUs associated with the individual SDLC interface.
9. Support of hunt groups for non-switched, fixed characteristics SDLC interfaces is desirable. If supported, the members of the hunt group consist of individual SDLC secondary stations for one or more such SDLC interfaces. In this case, a traditional hunting algorithm, as applies to PPSN X.25 interfaces (see the PPSNGR section 3.2) should be used.

5.3 Mapping XID Parameter Values

The XID (exchange station identification) frame is both a command sent by a primary SDLC station and a response from a secondary station. The XID command issued by a primary station requests that the secondary station identify itself via an XID response. The frame includes an optional information field to be used in identifying the sending station.

Although the use of XID is not restricted to switched SDLC interfaces, it is typically used immediately after a switched physical connection is established for the purpose of identifying or verifying the identity of the secondary station to the primary station, and optionally identifying the primary station to the secondary. Typically, a successful XID exchange is prerequisite to entering the normal response mode (or asynchronous balanced mode) on a newly established switched connection.

Over SDLC interfaces between two type 2.1 nodes, the XID exchange is used not only for identification but as part of the negotiation between the two peer nodes for primary and secondary station roles and other characteristics of the link. In this use of XID frames between type 2.1 nodes, ~~there~~ may be more than a single command and response to the exchange, as the capabilities of each node are established and the specific node roles and link characteristics are negotiated.

The SDLC PAD should be able to respond to XID commands with a valid response when acting on behalf of a secondary or type 2.1 node station. The PAD should also be able to initiate an XID command on behalf of a primary or a type 2.1 node station. Details of PAD support appear below.

1. Across the local SDLC interface, the SDLC PAD should be able to receive an XID command or send an XID response when directly connected to a primary station or type 2.1 node. Across the local SDLC interface, the SDLC PAD should be able to send an XID command or receive an XID response when directly connected to secondary

station(s) or a type 2.1 node. XID procedures and formats should be in conformance with IBM SDLC specifications^[6] in all cases and also in conformance with XID and XID3 specifications for type 2.1 nodes^[7] when the PAD is directly connected to a type 2.1 node.

2. Across a type A virtual circuit connection, the SDLC PAD should be able to send and receive QXID qualified Data packets in conformance with the IBM QLLC specifications,^[9] the general frame/packet mapping specifications of section 6.1.1, and the specific XID specifications of the following items.
3. For each PU configured on the local SDLC interface, the PAD should be configured with an XID information field value of up to 8 octets, to be used in populating XID frames when the virtual circuit connection is type B or the local SDLC interface is switched and serves a secondary station. When the PAD serves a primary station across the SDLC interface, this XID value is used to identify the (virtual) secondary station in response to an XID command from the primary when QLLC is not supported across the virtual circuit. When the PAD serves a secondary station across the SDLC interface, this XID value is used to identify the primary station emulated by the PAD in the XID command it sends immediately after the PAD establishes a dial-out SDLC connection to the secondary or immediately after a type B virtual circuit connection is established.
4. When the PAD populates the information field of an XID frame to be passed over the local SDLC interface, the value should be that received via a QXID packet from the remote end, if available, or the value configured for that local interface otherwise.
5. When the PAD populates the information field of a QXID packet to be passed to the remote end over the virtual circuit, the value should be that received via an XID frame across the local interface.
6. The PAD sends an XID response across the local SDLC interface to a primary station only in response to an XID command received from that primary station. This response directly follows receipt of the XID command when QLLC is not supported, but is delayed, until a valid QXID response from the remote end of the virtual circuit has been received by the PAD, when QLLC is supported.
7. The PAD sends an XID command across the local SDLC interface to a secondary station immediately following establishment of a dial-out connection to the DTE or following receipt of a QXID packet from the remote primary station.
8. The PAD sends a QXID packet over the virtual circuit to the remote end primary station only in response to a QXID packet command received from the remote end primary and only if QLLC is supported. This response is delayed until an XID command is sent to the local secondary station and a valid XID response is received.
9. The PAD sends a QXID packet over the virtual circuit to the remote end secondary station if QLLC is supported and the local primary or type 2.1 node station sends an XID command to the PAD over the local SDLC interface.

5.4 Virtual Circuit Setup/Clearing

The following apply to virtual circuits in general:

1. For each valid SDLC PAD interface category, Table 5-2 specifies the events or conditions which trigger the various virtual circuit establishment/clearing and SDLC link

mode transition actions that the PAD can take.

2. The SDLC PAD should clear each established SVC affected with cause "Out of order" (and not attempt to re-establish the call) and reset each PVC affected with cause "Out of order" if any of the following conditions is detected:
 - the physical level of the corresponding SDLC interface is declared unoperational, a switched circuit connection in the path is disconnected (see section 3.1), or an error condition is detected on the SDLC interface that is not recoverable at the link layer
 - the associated secondary station enters and remains in the disconnected mode for more than N_{dm} seconds, where N_{dm} can be set to any value from 1 to 30 seconds.
 - the associated secondary station does not acknowledge polls after the established number of retries (the retry maximum should be BOC-settable at least to values in the range 0-5)

5.4.1 Permanent Virtual Circuits

When the virtual circuit connecting two SDLC PADs or an SDLC PAD and an X.25 DTE or X.28 PAD is a PVC, there is no call establishment or clearing phase. PVCs are established, modified, and removed by service order. The service order is realized in terms of interface configuration option values that are coordinated across all links and network nodes in the virtual circuit path, including the SDLC PAD(s) at one or both ends.

1. The SDLC PAD should support PVC connections in conformance with sections 4.1 and 4.2.
2. When a PVC is deleted, the mapping relationship between the logical channel number that locally identified the PVC and the local SDLC station address with which it was associated is deleted. The affected local SDLC station address is not related to a logical channel number until a new virtual circuit serving the corresponding PU is established. See section 6.3 for additional details on PAD actions following deletion of a PVC.

5.4.2 Switched Virtual Circuits

5.4.2.1 All Virtual Calls

Switched virtual circuits (SVCs) or virtual calls are used to establish on-demand packet connections between DTEs served by two SDLC PADs or an SDLC PAD at one end and either an X.25 interface or X.28 PAD at the other end. SVCs are appropriate when terminal equipment needs to connect to host equipment at geographically dispersed locations at different times or if the duration of the connections through the PPSN are not long enough to justify a PVC.

Unlike PVCs, SVCs do require call establishment and call clearing procedures. Two basic alternatives are available for initiating these procedures. First, switched virtual connections can be initiated and cleared from the SDLC PAD in response to the SDLC station or entire link becoming active and becoming inactive, respectively. In this case, the SDLC uses stored call characteristics information and the automatic (auto) call service feature to originate or clear the call. A second alternative is for the customer to dynamically request the establishment or clearing of a virtual circuit involving one or more SDLC PADs. Under this latter alternative, procedures are provided to permit an authorized PPSN customer to request establishment and/or clearing of such virtual calls via a PPSN asynchronous (i.e., X.28) interface configured to support this function. Details specific to these two alternatives (PAD-initiated and DTE-initiated calls) are covered separately in two later subsections.

Although the ability to control virtual circuit establishment through the same terminal equipment that is to use the connection (i.e., the SDLC station equipment) is desirable, this requires the ability to communicate with terminals having diverse display characteristics, using higher layer SNA protocols. Requirements for support of this capability are under consideration.

Aside from the mechanism for initiating and clearing calls, there is one other major aspect of SVCs carrying SNA-compatible traffic that goes beyond the features addressed for normal X.25 and X.75 connections. That is the passing of logical link control (LLC) type and other SNA station details to the remote end as part of the call user data (CUD) field of the Call Request packet. NPSI and other SNA X.25 implementations that support QLLC need an indication of the type of LLC that applies to the virtual circuit connection. For an SVC, an indication of LLC type must be passed as the first part of the CUD field of the Call Request packet. To permit proper allocation of a virtual circuit to a destination station address at a grouped characteristics host interface serviced by an SDLC PAD (see section 5.1.4), it will also be necessary to pass some PU and LU characteristics information in the CUD.

1. The SDLC PAD should support SVC connections in conformance with sections 4.1 and 4.2.
2. Implementations should be consistent with the generalized chronological relationships among physical layer, link layer, and SVC establishment/clearing events diagrammed in Figure 5-2.
3. Implementations should conform to the specific relationships in time among the following (as illustrated in Figure 5-3):
 - establishment of dial-out connections, where applicable
 - establishment of the physical level of the PAD SDLC interface at the calling and/or called ends
 - completion of the XID exchange (when used) between the primary and relevant secondary station at the calling and/or called ends
 - entry into the normal response mode (or asynchronous balanced mode) on the link level of the PAD SDLC interface at the calling and/or called ends for the relevant secondary station
 - generation of the Call Request packet by the calling DTE or SDLC PAD on behalf of the calling DTE
 - acceptance of the Incoming Call or incoming Call Request packet by the called DTE or SDLC PAD on behalf of the called DTE
 - generation of the Call Accepted/Call Connected or Clear Request packet by the called DTE or SDLC PAD on behalf of the called DTE
4. The SDLC PAD should be capable of generating a Call Request packet that contains a CUD field. The presence of this field is mandatory for type A connections and is optional for type B connections. If the CUD field is present, it should consist of one, two, three, or four subfields. The following table provides an overview of the four possible subfields of the CUD:

Subfield(s) of the Call User Data (CUD) Field

Octet #	Subfield Contents	Presence
1	IBM QLLC Subfield	mand.
q		
q+p	SDLC PAD Subfield	opt.
q+p+r	Grouped Char. Port Subfield	opt.
q+p+r+s	User Data, If Any	opt.
t		

$1 \leq q \leq 12$ $p = 0$ or 1 $r, s = 0$ or $r = 1, s = 2$
 $q+p+r+s \leq t \leq 16$ (no fast select) or $q+p+r+s \leq t \leq 128$ (fast select)

- a. The first subfield of the CUD is mandatory and contains the protocol information specified by IBM for QLLC. The structure of this subfield is outlined by the following table:

The First CUD Subfield: Specified for IBM QLLC

Octet #	Bits of Octets							Presence
	8	7	6	5	4	3	2	
1	Protocol Identifier: P T X X L L C T							mand.
2	Format Field Identifier							opt.
3	ELLC Call Control Ind.							opt.
4	Reserved (binary all zeros)							if present
5	Connection							opt.
12	Identifier							

The first subfield of the CUD contains the protocol identifier, specifying the LLC type applying to the virtual circuit as its first octet. The coding of the first octet and the remainder of this first subfield should conform to sections 6.2.1.6 and 8 of Part II of the IBM NPSI specifications,^[9] and specifications of the PATH macro instruction and of the USRFILD operand of the X25OUFT macro instruction used by NPSI to populate the first part of the CUD.^[11] Specifically, the following apply to the first subfield:

- The first octet of the first subfield contains the protocol identifier.
- Bits 8 and 7 of the first octet (protocol type: "PT" in the above table) determine the use of the protocol identifier ("00" = CCITT X.29; "01" = use by network administrations; "10" = reserved for ISO use; "11" = used by higher layer protocol). IBM SNA X.25 DTEs assume they will always receive a CUD with bits 8 and 7 of the first octet set to "11".

- If PT = "11" then bit 2 of the first octet (next to least significant bit of "LLCT" semi-octet described below) is used to specify the type of connection (0 = SNA-to-non_SNA; 1 = SNA-to-SNA).
- The second semi-octet of the protocol identifier is used to distinguish the logical link control (LLC) type applying to the connection (the "LLCT" semi-octet of the above table).
- Optionally, octets may follow the protocol identifier. If present, the next octet is the Field Format Identifier, which defines the format of the following octets. Only the value x'01' is currently defined. If an IBM SNA X.25 DTE does not support the Field Format Identifier, it must accept and ignore up to 15 octets following the protocol identifier octet.
- The next octet of the first subfield is the ELLC Call Control Indicator, which is used to distinguish between initial connection requests (x'00') and connection recovery requests (x'01') for the ELLC protocol.
- The following octet of the first subfield is reserved and set to all zeros.
- The remaining eight octets of the first subfield, if present, contain the Connection Identifier, which not all SNA X.25 DTEs may support.
- The following table specifies the contents of the protocol identifier and any additional octets of the first subfield for various specific cases:

Protocol Identifier and Additional Octets of First Subfield for Specific Cases

Octet 1 in Hex	Protocol Type	LLC Type	Poss. Added Semi-octet(s)	Interpretation of Added Octets
C0	Native X.25	0	xxxxx	id # of remote PU [= 1 + id # of local PU] (00000 if not needed)
C2	PSH (for NIA)	2		
C3	QLLC	3	010000 (*)	
C4	GATE function	4		
C6	ELLC	6		
01	CCITT X.29	5	000000	
41	Network PAD	5	000000	
81	ISO PAD	5	000000	

* or 0100008dddwww000000 if IBM 3710 is cluster, where
ddd = upstream station address and wwwwww is the PU id # used for XIX
("00000" if a real PU type 2 is attached to the 3710)

The first octet of the CUD received or sent by an SDLC PAD should be x'C3' if QLLC is supported and x'C0' if QLLC is not supported over the VC. An SDLC PAD serving at the primary station end should also permit a first octet value of x'01' to be used when QLLC is not supported. Other protocol ID values are invalid for VCs involving an SDLC PAD.

- The second subfield of the CUD is mandatory whenever an SDLC PAD is or may be at the remote end of the connection, and is always present if the third subfield is present. The current version of this subfield consists of a single octet, encoded as follows:

Bits:	8	7	6	5	4	3	2	1
Octet	1	1	1	1	Vers.	Caller	Cont.	

The first semi-octet is set to all ones (x'F') as a flag to facilitate recognition of this subfield. The "Vers." bit designates which version of the encoding is being used. Currently, only the value "0" is used. The value "1" is reserved for possible use in designating alternate encoding schemes for the second and third subfields (to pass additional information needed to support multiple VCs per SDLC interface, for example). The two bits of "Caller" indicate the SDLC PAD role of the calling end:

- 00 = primary station (HPAD)
- 01 = secondary station (TPAD)
- 10 = negotiable role (SPAD)
- 11 = no SDLC PAD at calling end

The "Cont." bit indicates whether the third subfield is present (0 = not present; 1 = present).

- c. The third subfield of the CUD is present whenever the remote end is or may be a non-switched, grouped characteristics SDLC PAD interface. This subfield is used to pass PU and LU characteristics information that is needed when the destination DTE is served by a grouped characteristics SDLC interface on a PPSN PAD. This information is used by the SDLC PAD to assign an incoming call to an appropriate station address on a grouped characteristics host end interface or interface hunt group (see section 5.1.4). This subfield, when present, consists of 2 octets:

Bits:	8	7	6	5	4	3	2	1
2nd Octet	PUa	# of LUs; Formatted Commands for SSCP Session						
3rd Octet	PUB	# of LUs; Character-Coded Commands for SSCP Session						

The PU Type is encoded in the PUa and PUB bits. If PUa = 0 and PUB = 0, the PU is type 1. If PUa = 1 and PUB = 0, the PU is type 2. Interpretations for cases with PUB = 1 are reserved. The remaining (least significant) seven bits of the second and third octets encode the number of LUs of two types associated with the PU. The values are coded in binary with bit 1 being least significant and bit 7 being most significant. The second octet corresponds to the case of SSCPFM=FSS for the VTAM LU statement. The third octet corresponds to the SSCPFM=USSCS case. The sum of the second and third octet values should equal the total number of LUs associated with the PU.

- d. The fourth subfield of the CUD is optional, and contains any discretionary user data to be appended to the CUD for user application purposes. The total of the four subfields cannot exceed 128 octets if the Fast Select facility applies or 16 octets, otherwise.
5. When an SVC is cleared, the mapping relationship between the logical channel number that locally identified the SVC and the local SDLC station address with which it was associated is deleted. The affected local SDLC station address is not related to a logical channel number until a new virtual circuit serving the corresponding PU is established. See section 6.3 for additional details on PAD actions following the clearing of

an SVC.

6. Each actual or emulated secondary station on each terminal-end and host-end SDLC interface should be capable of being configured for either PAD-initiated (auto call service) or DTE-initiated calls (the DTE being a terminal acting on behalf of the interface from which the Call Request will be actually initiated). See the following two subsections, respectively, for details on these two cases. A given secondary station can only be configured for one of PVC, PAD-initiated SVC, or DTE-initiated SVC service. Within the constraints of the total number of PVC and SVC logical channels available and the number of secondary stations configured, the overall SDLC interface should be capable of simultaneously supporting PVCs, PAD-initiated SVCs, and/or DTE-initiated SVCs.
7. The PAD can initiate a virtual call only if the corresponding configured secondary station at the local SDLC interface is in the NRM, NRME, ABM, or ABME condition. If the secondary station leaves any of these link modes for more than a threshold period (N_{dm} , defined below), the call should be cleared. A Call Request incoming to a PAD SDLC station can be completed only if the station on the local receiving SDLC interface is in the NRM, NRME, ABM, or ABME condition. For non-switched interfaces, if the configured secondary station is not in one of these link modes, the call request should be cleared with cause "number busy" (or "out of order" if the interface or station is declared inoperable at the PAD). It is assumed for non-switched interfaces at the terminal end, that local PAD polling activities (section 3.2.2.3) assure that the secondary station is in one of the above four link modes if the station is prepared to enter that mode.

A call request incoming to an SDLC PAD that requires a switched connection to reach the destination DTE should initiate an attempt to establish that switched connection and subsequently establish the link layer for that connection. The call request should not be accepted until and unless the switched connection is established and the associated link layer can be placed in the NRM(E) or ABM(E) condition, as appropriate.

8. It is desirable for the SDLC PAD that issues the Call Request to be able to automatically clear an established SVC according to one or more of three time-based criteria. This capability is labeled timed auto call clearing. The three criteria are:
 - at a specified time (hh:mm:ss) within the next 24-hour period
 - a specified duration (hh:mm:ss) after the Call Request was sent
 - a specified duration (hh:mm:ss) during which no unqualified Data packet was sent or received over the virtual circuit

If this capability is supported, it should be possible to select the option and the applicable time value on an interface-specific basis. When the call is cleared in this manner, the cause should be "DTE originated" (00000000).

9. If an SVC is cleared with any cause in the following list, then the SDLC PAD that placed the Call Request should increment a retry counter and re-establish the call (following the applicable T_d period specified in section 6.3) if the retry counter has not reached the retry limit configured:
 - Remote procedure error
 - Local procedure error

- Network congestion

Once the retry limit is exceeded, the counter is reset and no further call establishment attempts are made by the PAD without operator intervention (for the auto call feature) or explicit call establishment request via an asynchronous control interface. The re-attempt limit should be settable at least to values in the range 0-5. After a normal "DTE originated" clearing (cause and diagnostic codes of decimal value 0), the retry counter is reset.

10. If an SDLC PAD receives a valid incoming call request after having itself sent a call request with no response from the remote end, both of which apply to the same secondary station address, the call collision is handled as follows. A PAD interface configured to support a local secondary or negotiable secondary/primary end clears the incoming request with cause "Number busy" and diagnostic code decimal 72. A PAD interface configured to support a local primary station accepts the incoming request if it is otherwise valid and assumes its own call request will be cleared by the remote end. If its request is not cleared by the remote end within 200 seconds (or is accepted), the PAD issues its own clear request to free up that logical channel number.

5.4.2.2 PAD-Initiated Calls

Configuring an SDLC secondary station for PAD-initiated calls is appropriate when calls from that station are always made to the same physical destination and with the same call characteristics. The PAD uses call information stored in memory to place the call automatically, in response to a transition in the status of the physical interface or the link station, or certain types of call clearing. For this SVC alternative, the auto call service provided by the PAD reduces to a minimum the action necessary on the part of the customer in setting up the virtual circuit connection, at the expense of reduced flexibility. The PAD-initiated call differs from PVC service in that the call is active and network logical channel resources are reserved only for a certain duration, defined by configuration parameters and changes in interface or station status.

1. Support of auto call service by the SDLC PAD should be provided on non-switched interfaces. This service permits PAD-initiated calls to be made. If the SDLC interface is configured for auto call service, the PAD should be capable of automatically establishing and clearing SVCs, in response to transitions in the status of the physical level of the SDLC interface and the link level mode of each (actual or emulated) secondary SDLC station. Certain types of call clearing can also trigger a PAD-initiated call, as specified below. The Call Request and Clear Request packets to establish and clear the SVC, respectively, will be generated based on service data configured for the interface and each SDLC secondary station address to have the auto call service. This service data will consist of an image of the Call Request and Clear Request packets to be used, including addresses, all per-call facilities to be included (checked for consistency with applicable subscription facilities), and user data. If timed auto call clearing is not active, a call created via the auto call feature will be cleared in accordance with item 2 in section 5.4.
2. For each configured (actual or emulated) secondary station on the interface, the PAD will initiate auto call service for a virtual circuit if there is no established virtual circuit and the secondary station has just entered the NRM, NRME, ABM, or ABME mode on the link.
3. The PAD will also initiate auto call service for a virtual circuit if that circuit was cleared with those causes specified in item 9 of section 5.4.2.1 and the retry limit has not been exceeded.

5.4.2.3 DTE-Initiated Calls

Configuring an SDLC secondary station for DTE-initiated calls is appropriate when calls from that station are likely to be made to different physical destinations and/or with varied call characteristics. Unlike the auto call feature, this service option provides full call establishment flexibility, but the customer must explicitly specify the SVC to be established. The customer also specifies the criterion to be used for clearing the virtual circuit. In addition, the customer may subsequently request clearing of the SVC in real time.

An authorized customer uses a specially configured asynchronous (X.28) control interface to request the establishment of SVCs originating from SDLC interfaces served by the same AC or PS that contains the control interface (see Figure 1-3). An authorized customer can set up SVCs from multiple stations and SDLC interfaces on the AC or PS from a single control interface. The control interface can also be used by the customer to receive status and clearing messages for all SVCs for which it is authorized.

Authorized control interfaces are asynchronous interfaces configured to support an extension of the command and display protocol used on normal asynchronous interfaces, as specified in section 3.3 of the PPSNGR.^[3] Non-switched asynchronous interfaces may be configured as SDLC PAD control interfaces. In addition, specially configured dial-in asynchronous interfaces may be used to request that a secure dial-back asynchronous interface be established as a control interface.

1. The AC should support non-switched asynchronous interfaces that can be configured for the SDLC interface control function. Such interfaces would support an extension of the X.28 protocol specified in section 3.3 of the PPSNGR. The extension would permit the interface to be used to establish and monitor switched virtual circuit connections that originate from SDLC interfaces on the same AC. It is desirable that all asynchronous interfaces on an AC with SDLC PAD function be capable of being configured for the SDLC interface control function. It is desirable for the PS to support such interfaces if the PS also supports the SDLC function.
2. It is desirable for the AC and PS to support dial-in asynchronous interfaces that can be used to request establishment of a dial-out asynchronous control interface from the same AC or PS. Once established and validated, and until inactivated, the dial-out control interface would operate like a non-switched control interface. See section 5.7.1 for relevant security procedures used for dial access.
3. Each authorized user on a non-switched control interface should be configured with a list of all of the SDLC interfaces or SDLC interface/station address combinations for which control by that interface was authorized by the owner of the respective SDLC interface. Each station, interface, or group thereof is identified through its PPSN DTE or hunt group X.121 address. The list should be able to accommodate at least five different X.121 addresses (full/individual and/or truncated/group addresses, as described in Appendix C). At least two authorized users (each identified by an authorization code) per non-switched interface should be supported. On each dial-up control interface, at least 100 authorized users should be supported. (See section 5.7.1 and Appendix C for further details.)
4. An asynchronous control interface should operate in accordance with section 3.3 of the PPSNGR when its SDLC interface control function is not being exercised. An extension of the normal asynchronous protocol specified in section 3.3 of the PPSNGR is used for the control function. An integral part of this protocol extension is a unique escape from the regular asynchronous protocol, which clearly indicates when the control mode is active. Appendix C specifies the control extension protocol that applies for the control function.

5. An asynchronous control interface should be able to receive SVC status and clearing messages for the same set of SDLC interfaces or SDLC interface/station address combinations for which control is authorized. Messages would be received while the interface control session is active for those SVCs initiated from that interface during the same session or explicitly requested during that session. The format for requesting message display is also specified in Appendix C. Appendix D specifies the format to be used in displaying various call status and clearing messages.
6. Requirements to support screen-oriented input command and message display formats for either asynchronous control interfaces or SDLC interfaces are under consideration. However, if menu-driven screen formats were offered on asynchronous control interfaces or directly to originating SNA SDLC devices, it is desirable that a consistent format be adopted. Appendix E provides recommended formats for command/request and display screens that could be used in place of the line-oriented formats of Appendixes C and D if the more user-friendly screen formats were supported, in addition, as an interface configuration option.

5.5 Virtual Circuit Profiles

To facilitate the establishment of SVCs by PPSN customers and to reduce memory requirements associated with specification of PAD-initiated SVCs using auto call service, the concept of virtual circuit (VC) profiles is introduced. A VC profile is a full or partial specification of an SVC that can be used in constructing and executing a Call Request (and, optionally, a subsequent Clear Request) for an SVC. The profile is stored in network memory, accessible to the SDLC PAD that will be responsible for establishing and clearing the SVC. A PPSN customer or auto call specification refers to a VC profile with a decimal profile number.

The profile number may be unique across the PPSN, within the AC or PS containing the SDLC PAD function, or unique only to the X.121 address of the control interface specifying the profile number. The decimal value of the profile number determines which of these degrees of uniqueness applies. A one-digit VC profile number (i.e., values 0-9) can be used only in a control interface SVC establishment request and is unique only to the X.121 address specifying the profile number. Profile numbers in the range 10-19 are reserved to be uniform across all PPSNs, as specified in this document and future issues of the PPSNGR (see Table 5-3). Profile numbers in the range 20-29 are reserved to be uniform across all ACs and PSs within a PPSN, as set by the responsible BOC. The remaining two-digit profile numbers (30-99) are unique within the AC or PS containing the auto call specification or the control interface across which the number is passed.

One-digit profiles would be used by individual customers to store the details of up to 10 SVCs commonly initiated over the respective control interface. Two-digit profiles could be referenced by any control interface and/or for any calling address on the AC or PS storing the profile. These would be useful for profiles shared by multiple users and for auto call service specifications.

1. Each AC and PS supporting the SDLC PAD function should support one-digit and two-digit VC profiles. An AC and PS should be able to store at least 5 profiles in the range 0-9 for each SDLC PAD control interface configured on the AC or PS. The AC and PS should be able to store profiles that are maintained on a uniform basis throughout the network for each value in the range 10-29. The AC and PS should be able to store at least 25 profiles (which need not be related to profiles with the same number in other ACs and PSs) in the range 30-99.

2. When a VC profile is specified, it is used by the AC or PS to construct and issue Call Request and Clear Request packets for the associated SDLC PAD SVC. A VC profile number may be specified by a user over an authorized control interface in a SDLC PAD SVC establishment request (see Appendix C) or (in the case of two-digit profiles) be referenced by one or more auto call service specifications (to reduce storage requirements and facilitate network data base maintenance). Any specification stored in a VC profile is supplemented or overridden by specifications that explicitly accompany a reference to that VC profile in an SVC establishment request or auto call service specification.
3. A partial VC profile does not contain one or more elements of a full profile which are expected to vary from call to call (e.g., calling and/or called address). The missing information of a partial profile must be provided by the command request or auto call specification details that accompany the reference to the profile. A full VC profile contains the following specifications:
 - calling and called DTE addresses
 - which user facilities appear in the facility field and the coding details for each facility that appears
 - the contents, if present, of the call user data field (permitting flexibility to include hexadecimal semi-octet values) up to 16 octets (128 octets if fast select is active)
 - if clearing of the call is pre-specified, the time at which the call is to be cleared (see Appendix C for the three alternative ways for specifying clearing time)
 - the contents, if present, of the user data field of the Clear Request packet up to 16 octets or 128 octets if fast select is active

5.6 Data Transfer Phase Functions

At all times for a PVC and during the data transfer phase of an SVC, the SDLC PAD performs several functions to assure adequate network performance (aggregate and per-VC), mediate between incompatible frame and packet sizes, and process any interrupts or resets received. Local polling (section 3.2.2.3), processing speed, and the buffers provided should permit the performance objectives specified in section 9 of the PPSNGR to be achieved. When the information fields of I frames on the local SDLC interface and the user data field of virtual circuit Data packets are size incompatible, the SDLC PADs perform segmentation and recombination functions, as specified in section 6.1, to achieve compatibility. When Interrupt and Reset packets are received, they are handled in accordance with the virtual circuit support specifications of sections 4 and 5.4, SDLC/VC mapping rules of section 6, and the specifications in Recommendation X.25 and section 3.2 of the PPSNGR.

5.7 Security Features

Two categories of security features are appropriate when SNA applications are supported across a public packet switched network: (a) access control and (b) disconnection protection. It is important to control access to both host applications and control interfaces beyond the level of security afforded by normal virtual circuit procedures. The need to carefully establish authorization when a control interface is used to set up virtual circuits on behalf of other interfaces is obvious, especially when control is requested via a public dial-in port. The need for enhanced security measures when a host application is accessed via a PPSN virtual circuit connection

results because the presence of a virtual circuit may distort the host's perception of the end user. Specifically, if the host end SDLC interface to the SDLC PAD is non-switched, a host will perceive the user to be "dedicated" and treat security accordingly, even if the overall connection was established on a switched basis (i.e., via an SVC and perhaps via dial-out on the remote end of the SVC).

Since diverse public users can theoretically access an SDLC interface supported by a PPSN PAD, which a host perceives as being "dedicated," there is also a need for the network to provide special protection if the connection is broken. If a switched virtual connection is cleared, it is important that the host be made aware of this fact before any other user is allowed to establish a virtual connection to the same host end PU as in the original (lost) virtual circuit connection. Without such protection, another user might be able to access the same host application without the host being aware of the change.

The following two subsections deal with additional security mechanisms appropriate when virtual circuits are established in support of SNA applications.

5.7.1 Access Control

Access control consists of several mechanisms to assure that a user establishing a virtual circuit connection to an SNA host is authorized to access the host port. These mechanisms are available since access may appear to be dedicated to the host even if a switched connection is actually involved. The PPSN makes use of NUI, XID, closed user group (CUG), authorization code, and secure dial-back mechanisms to provide adequate access control protection in these cases.

1. The SDLC PAD should support authorization code verification to establish that asynchronous control interface users are authorized (see Appendix C for details).
2. The SDLC PAD should support link level station identification over the local SDLC interface and remotely over virtual circuit connections supporting QLLC, as specified in section 5.3.
3. All non-switched SDLC interfaces for which the SDLC PAD is serving a primary station across the local interface should be capable of being configured to accept and/or initiate only calls within a CUG. If the CUG restriction is adopted for such an interface, it would prevent any DTE address that does not belong to a valid CUG associated with the host port from accessing that port interface. NUI and NUI Override may be used by dial-in users of SDLC control interfaces for billing and terminal profile purposes. However, the CUG interlock code that applies for an SVC is that of the calling SDLC interface DTE address and not that of the control interface specifying the SVC or any interlock code associated with a NUI provided by that control interface user.
4. All dial-in asynchronous interfaces on SDLC PADs configured to support the SDLC PAD control interface function should be capable of secure dial-back on the same or another interface on the PAD (see also Appendixes C, D, and E). After verifying that the dial-in control request is authorized, the PAD will break the dial-in connection and initiate a dial-out to the DTE address associated in memory with the valid authorization code provided by the dial-in user. Once the secure dial-back connection is successfully established and the correct authorization code again provided, this dial-out connection may be treated in the same fashion as a dedicated asynchronous control interface, as long as the connection is not broken (see section 3.1 for physical layer interruption time criteria).

5. All switched SDLC interfaces supported by SDLC PADs should be constrained to be dial-out to secondary SDLC stations across the local interface. Support of other switched SDLC interface options is under consideration.

5.7.2 Disconnection Protection

If an SVC involving an SDLC PAD at the primary station end is cleared other than at the instigation of a host application, that host application may have no indication of the interruption of the connection or of a possible breach in access security at the SNA session level. Without special security precautions, it might be possible for another authorized (but different) user to gain access to the same application session via a new SVC without that application being aware of a change in secondary station end LUs (i.e., the "session tailgating" problem). Thus, if SDLC PADs are used to locally terminate SVCs at the host end, additional security precautions will likely be necessary. Some such precautions may require the SDLC PAD locally connected to the host FEP to intervene at a protocol level above the link level (i.e., SNA messages to inform the application of the interruption at the session level).

1. Section 6 addresses actions taken by the SDLC PAD at the SDLC link level when an SVC is cleared. What action, if any, an SDLC PAD locally connected to a primary station via a non-switched SDLC interface should take to notify the host end at higher layers of the SNA protocol stack if the SVC terminating in that interface is cleared is currently under consideration. Specific PAD requirements for higher layer SNA session notification, in response to SVC clearing, may be proposed if they are found to be justified and workable. [For example, transmission to the SSCP of an exception request or negative response with an appropriate sense data value.]

In the absence of specific network security mechanisms at the LU session level, some PPSN customers may choose to:

- restrict access to non-switched SDLC interfaces at the host end to PVCs; or
- configure all (even non-switched) SDLC interfaces as switched in the FEP (but not necessarily in the SDLC PAD) if they are accessible via SDLC PAD SVCs.

Figure 5-1. Illustrating SNA Support of Simultaneous Access to Different Hosts


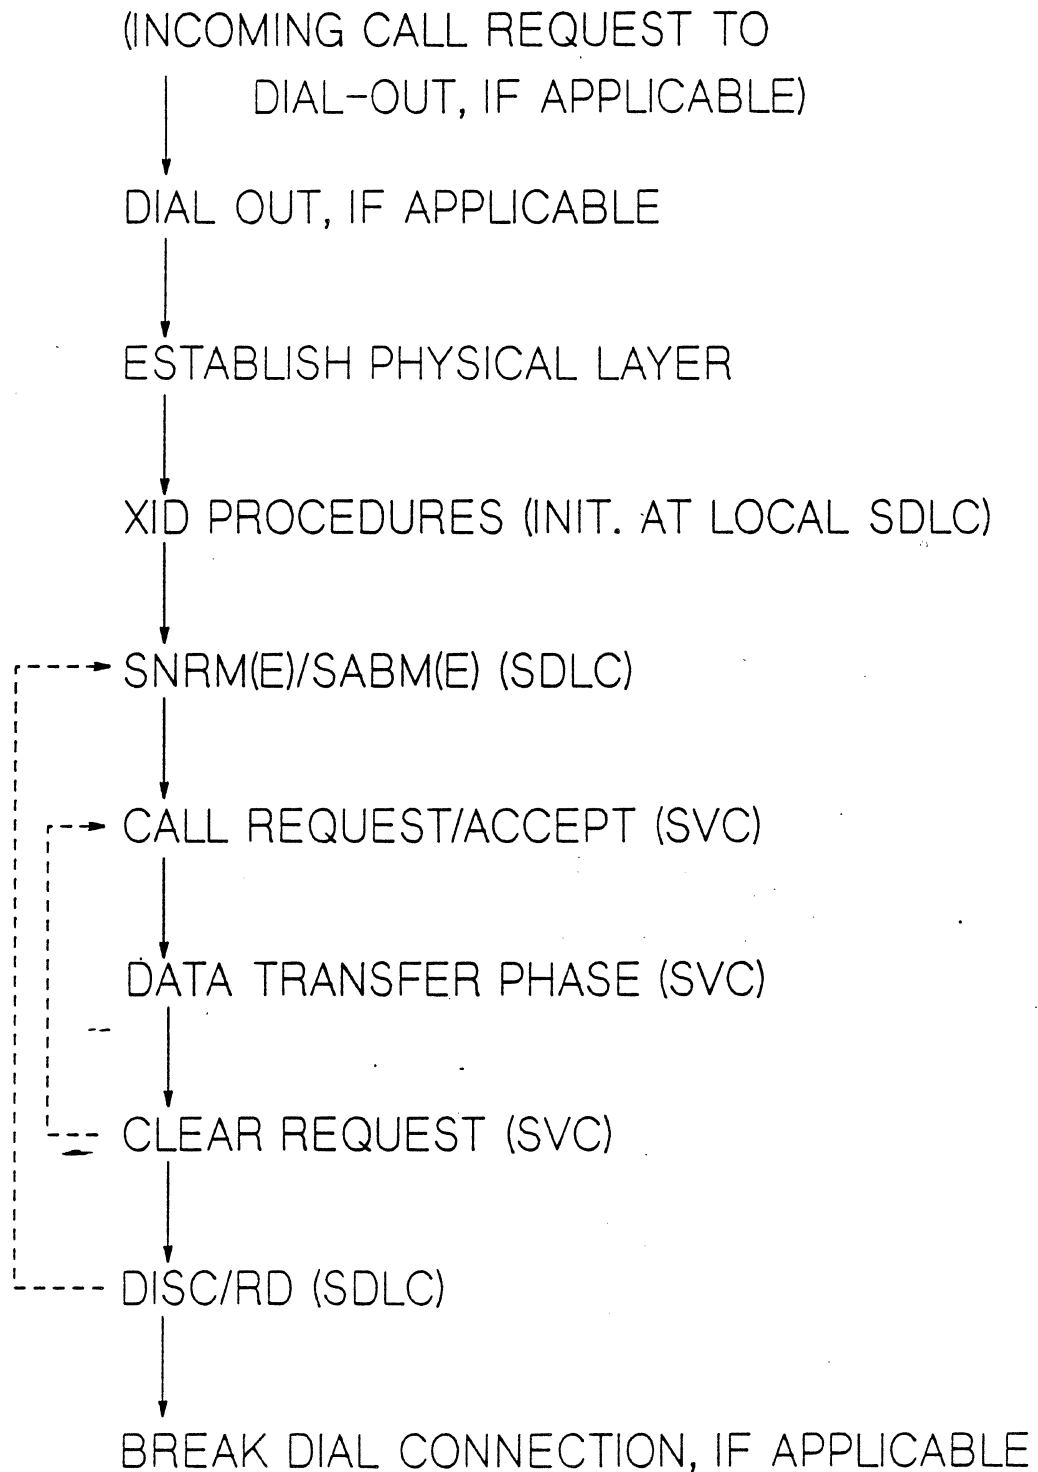


Figure 5-2. Normal Chronological Relationship Among Physical, Link, and SVC Level Events



PAD-INITIATED CALLS: CALLING DTE END

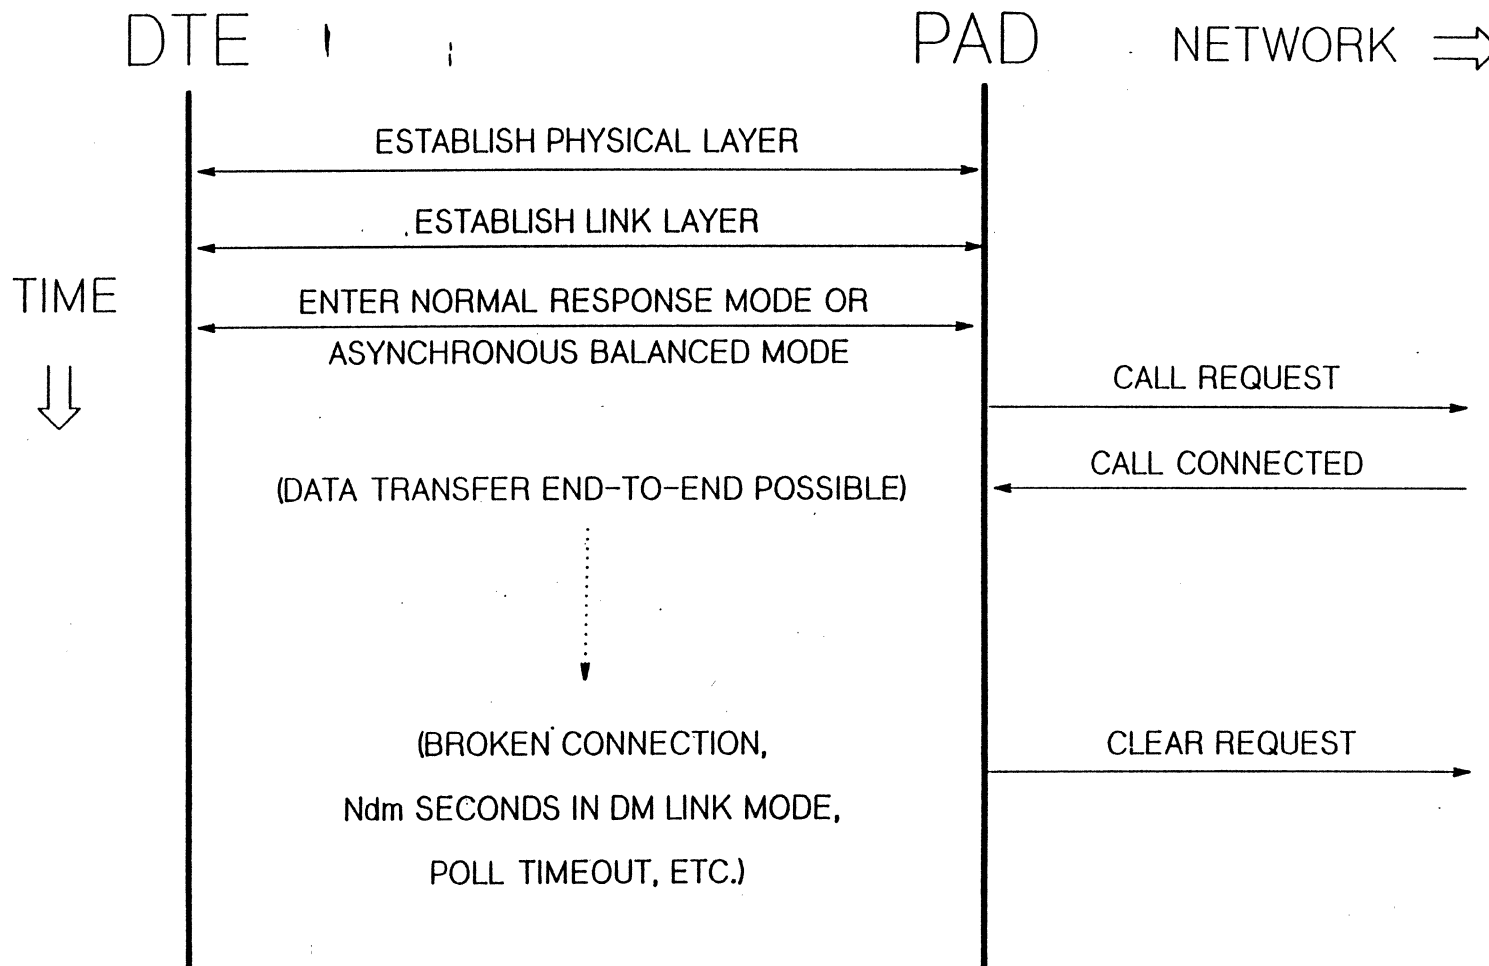


Figure 5-3. Time Relationships Among Link Level and SDLC Establishment Events

DTE-ORIGINATED CALLS: AT CALLING DTE END

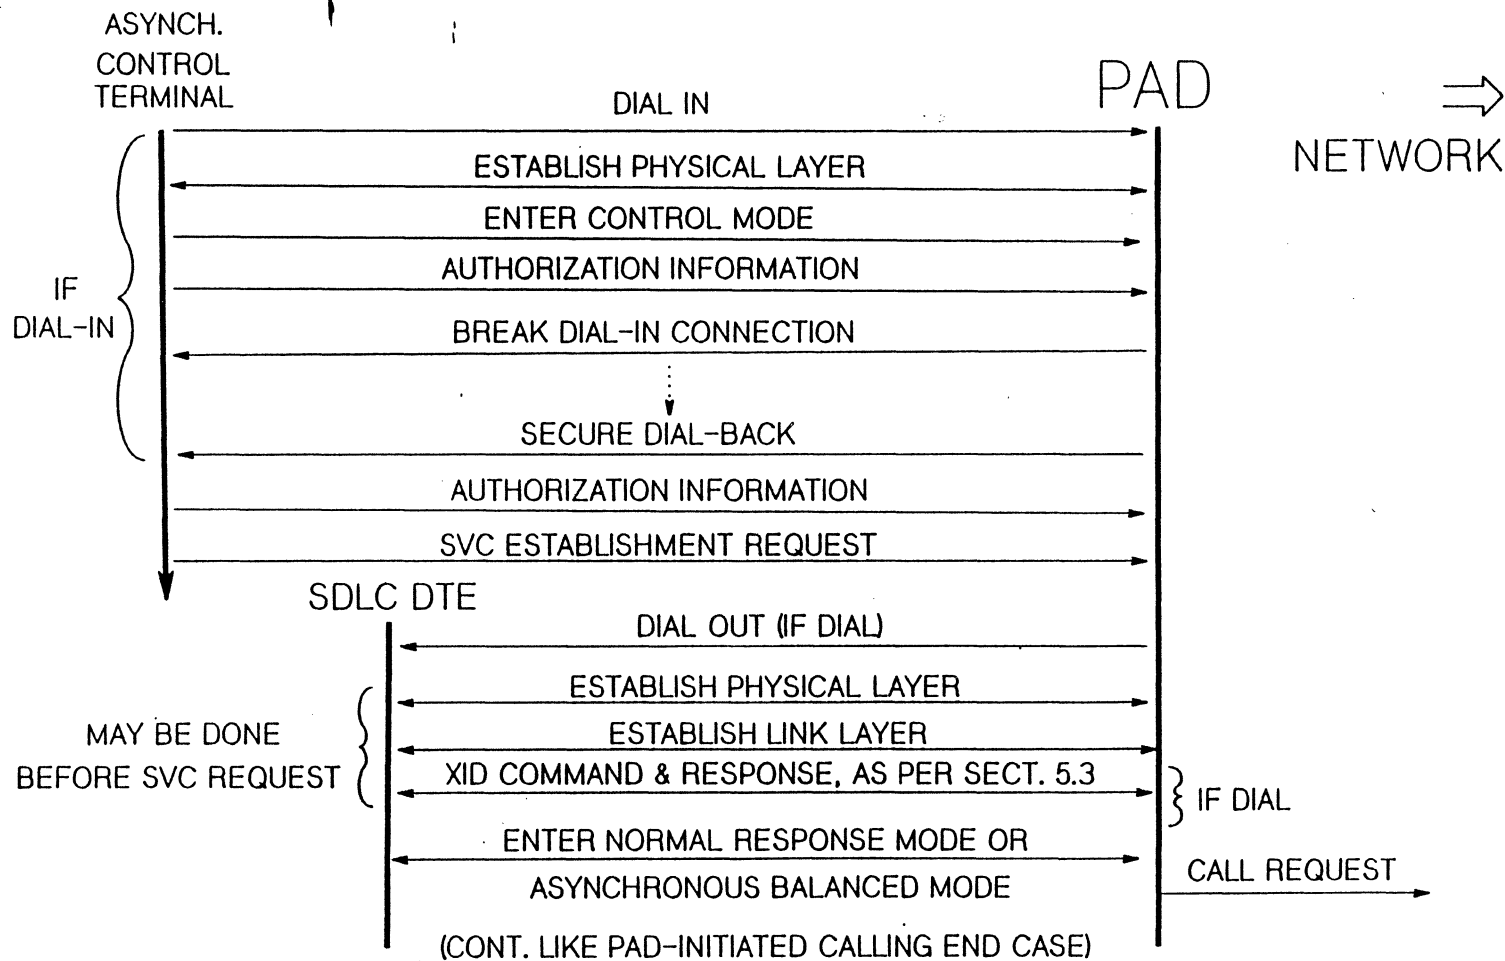
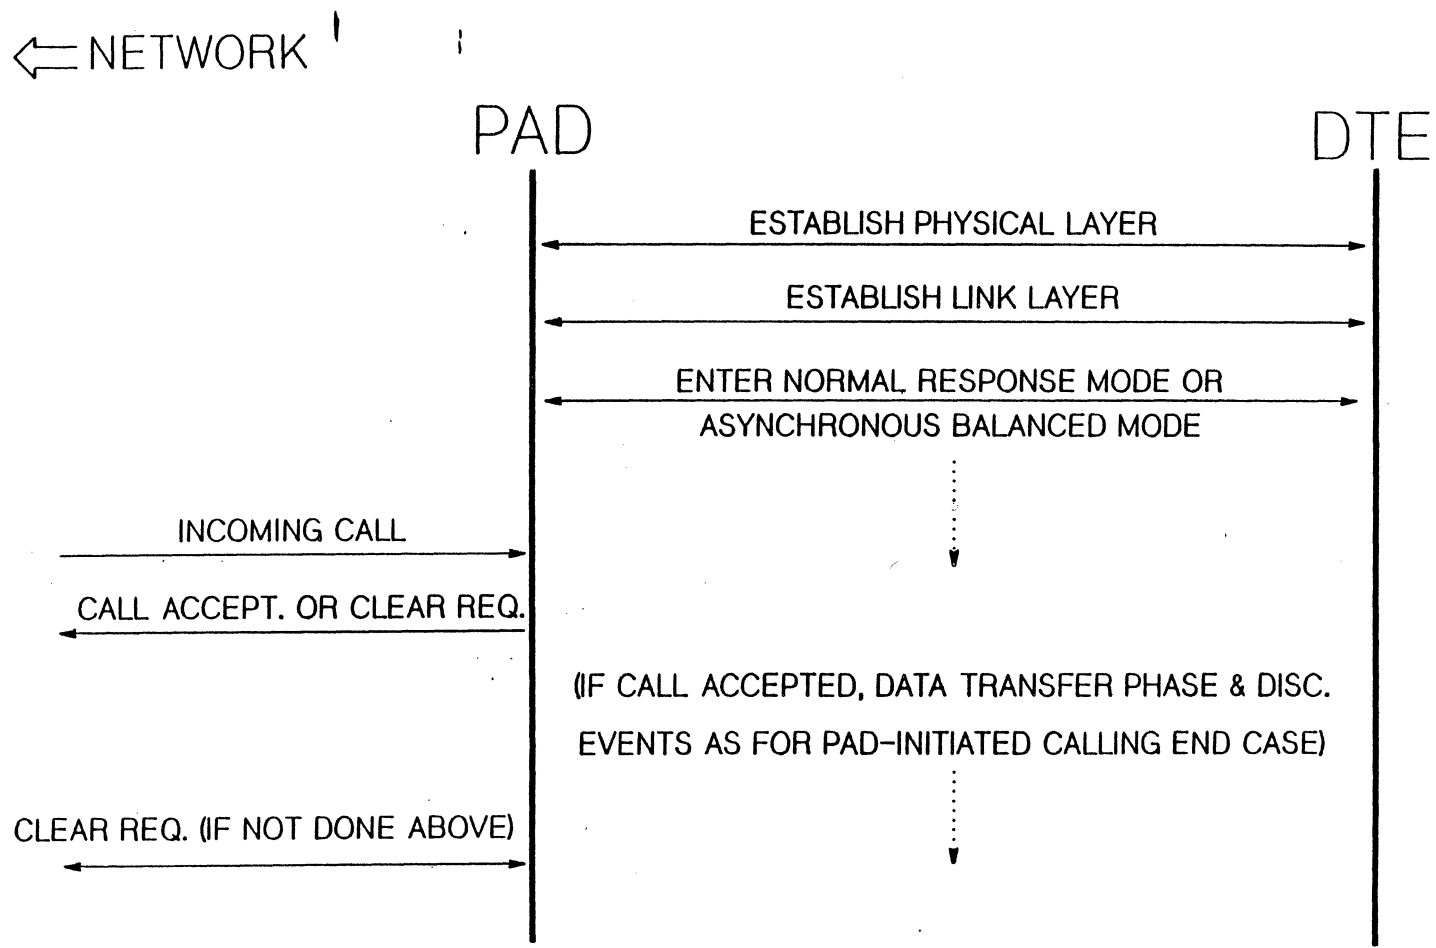


Figure 5-3. Time Relationships Among Link Level and SDLC Establishment Events (Cont.)

Figure 5-3. Time Relationships Among Link Level and SDLC Establishment Events (Cont.)

CALLED DTE END: NON-SWITCHED SDLC INTF.



CALLED DTE END: SWITCHED SDLC INTF.

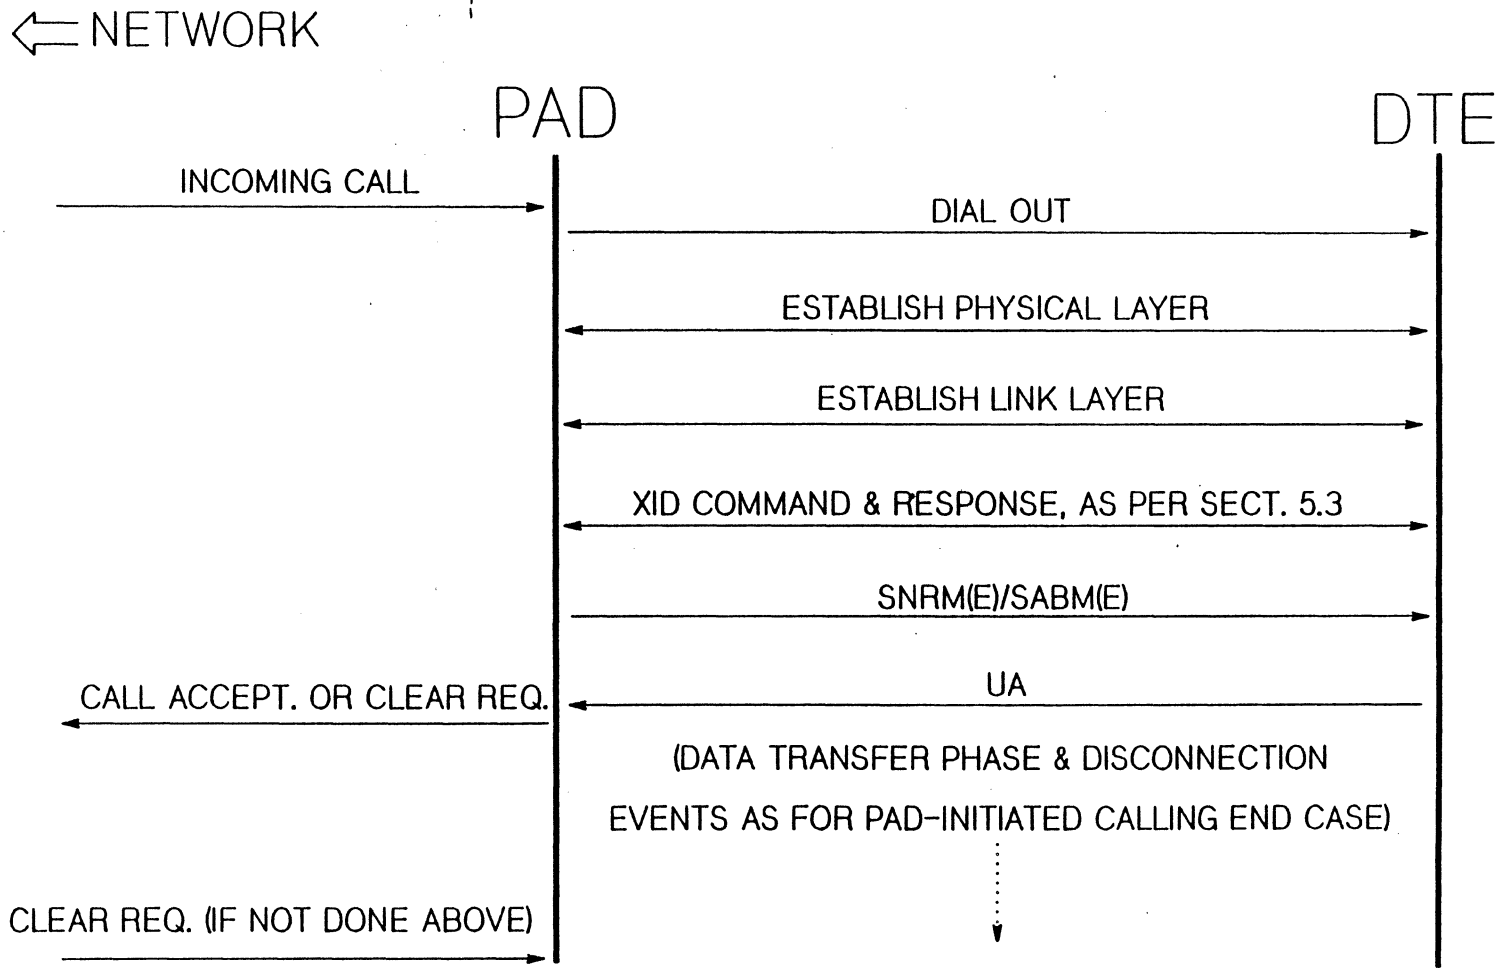


Figure 5-3. Time Relationships Among Link Level and SDLC Establishment Events (Cont.)

Table 5-1. Valid and Invalid SDLC PAD Interface Configuration Combinations

Local Station Type Served	PAD-Initiated SVC (Auto Call)			PVC	DTE-Initiated SVC		
	Switched	Fixed Char.	Grouped Char.		Switched	Fixed Char.	Grouped Char.
Secondary	NO	OK	NO	OK	OK	OK	NO
Primary	NO	OK	NO	OK	NO	OK	OK
Negotiable	NO	OK	NO	OK	NO	OK	OK

Table 5-2. Triggers for VC Establishment/Clearing and Link Mode Transition Actions

Local Station Type Served*	PAD-Initiated SVC, Fixed Char. (1)	PVC (2)	DTE-Initiated SVC		
			Switched (3)	Fixed Char. (4)	Grouped Char. (5)
Secondary (A)	<p>CallReq: At transition from dm to NRM/ABM, if sta. becomes active (sect. 3.2.2.3), or as part of ClrCR. Clr!: Specified time criterion, N_{dm} or poll timeout, or nonoper. SM: At PwrOn or if SVC established and according to polling procedures of sect. 3.2.2.3. DISC: If nonoper., as optional precursor to SM, or after CLR! of SVC.</p>	<p>SM: At PwrOn or if becomes oper. DISC: If PVC deleted or nonoper. or as optional precursor to SM.</p>	<p>CallReq: Same as A4. Clr!: Same as A4. MDC: When corresponding (calling/called) address specified by control intf. CallReq is on PSTN. BDC: When associated SVC is cleared. SM: Same as A1. DISC: Same as A1.</p>	<p>CallReq: At async. control intf. request or as part of ClrCR. Clr!: Time criterion or explicit clear req. specified by control intf., N_{dm} or poll timeout, or nonoper. SM: Same as A1. DISC: Same as A1.</p>	NA
Primary (B)	<p>CallReq: Same as A1. Clr!: Same as A1. DM: If in dm as result of DISC command or if nonoper. RD: If nonoper. while in NRM/ABM or Clr! or ClrCR of SVC.</p>	<p>DM: Same as B1. RD: If PVC deleted or nonoper. while in NRM/ABM.</p>	NA	<p>CallReq: Same as A4. Clr!: Same as A4. DM: Same as B1. RD: Same as B1.</p>	<p>CallReq: NA (call initiation from grouped char. intf. not appropriate) Clr!: N_{dm} or poll timeout or nonoper. DM: Same as B1. RD: Same as B1.</p>

(Legend on following page)

Table 5.2. Triggers for VC Establishment/Clearing and Link Mode Transition Actions (Cont.)**Legend**

*	SDLC interfaces configured to support DTEs that can negotiate a primary or secondary role adopt the station role that is actually negotiated through the XID process for Type 2.1 nodes.
CallReq	Call request generated (for SVCs only)
ClrCR	Clear that is followed by an attempt to re-establish SVC (sect. 5.4.2.1)
Clr!	Clear request not followed by call request retry (sect. 5.4.2.1)
SM	SNRM(E) or SABM(E), optionally preceded by DISC (by PAD at secondary only)
DISC	DISC command issued (by PAD at secondary only)
RD	RD issued (by PAD at primary only)
DM	DM issued (by PAD at primary only)
dm	NDM or ADM (disconnected) link modes
NRM/ABM	Normal Response Mode or Asynchronous Balanced Mode
nonoper.	a station or associated physical facility, link interface or VC declared nonoperable (fatal error; power off; "out of order" reset)
oper.	station and associated physical facility, link, and VC operable
PwrOn	station power on detected or assumed (e.g., return of carrier, etc.)
MDC	Make (establish) dial connection (switched only)
BDC	Break dial connection (switched only)
NA	Not applicable or invalid.

Table 5-3. Two-Digit (Partial) VC Profiles With a Uniform Definition Across PPSNs

Prof. No.	Description	CUD Field in Hex*	Clear Time	Rev. Chg?	Pkt Size	Wind. Size	Thruput Class**
10	RC, Interactive, QLLC	C3010000	_01:00:00	Yes	128	2	2400
11	RC, File Trans., QLLC	C3010000	_00:10:00	Yes	128	2	9600
12	No RC, Interactive, QLLC	C3010000	_01:00:00	No	128	2	2400
13	No RC, File Trans., QLLC	C3010000	_00:10:00	No	128	2	9600
14	RC, Interac., Non-SNA X.25	C0xxxxxy	_01:00:00	Yes	128	2	2400
15	RC, File Trans., Non-SNA X.25	C0xxxxxy	_00:10:00	Yes	128	2	9600
16	No RC, Interac., X.25	C0xxxxxy	_01:00:00	No	128	2	2400
17	No RC, File Trans., X.25	C0xxxxxy	_00:10:00	No	128	2	9600
18	RC, X.28/X.29 PAD	01000000	_01:00:00	Yes	128	2	1200
19	No RC, X.28/X.29 PAD	01000000	_01:00:00	No	128	2	1200

RC =reverse charged

* xxxxx = 5-digit identity of called PU (00000 if not needed)
 y = 0 pad digit to make even number of semi-octets, if needed

** If value is greater than interface line speed, then line speed value.

BLANK PAGE

6. SDLC/VC Mappings and State Changes

6.1 Mapping SDLC Frames to/from Packets

6.1.1 Default Case: Support of QLLC

When QLLC is supported by the SDLC PAD over the virtual circuit connection, the requirements of section 3.2.2.2 apply. The format and use of the supported QLLC qualified Data packets should conform to the IBM QLLC reference document cited earlier^[9] (in particular, section 8 in Part II of that document). For implementations of the mapping between QLLC qualified Data packets on the VC connection and frames from and to the corresponding local SDLC interface station address, the following procedural details also apply:

1. When an SDLC PAD receives an unqualified (data transfer mode, Q bit = 0) Data packet, the user data of the packet is transparently transferred to the Information Field of an I-Frame that the PAD passes over the local SDLC interface with the SDLC station address corresponding to the VC (i.e., LCN) of the received Data packet. However, when the received unqualified Data packet has the M bit set to 1, the PAD appends the contents of the user data field to the Information Field of the I-Frame being prepared, but does not yet transmit the I-Frame. The I-Frame is transmitted only after the user data field of a Data packet with the M bit set to 0 has been appended to complete the I-Frame Information Field. It is assumed that complete packet sequences (a contiguous sequence of zero or more Data packets with M bit set to 1 followed by a final packet with M bit set to 0) will not result in an I-field Information Field which exceeds the maximum frame size. If the PAD receives a Data packet sequence that exceeds the maximum frame size on the corresponding local SDLC interface, it discards the current I-Frame being accumulated and either clears (SVC) or resets (PVC) the corresponding virtual circuit (cause code "DTE originated" and diagnostic code decimal 161).

When an SDLC PAD receives a Clear Request or Call Connected packet with a valid user data field, the associated user data is transparently transferred to the information field of an I-frame that the PAD passes over the local SDLC interface with the SDLC station address corresponding to the VC of the received packet. The same type of mapping is performed with the fourth subfield of the CUD, if present, in a received Call Request or Incoming Call packet (see item 3 in section 5.4.2.1) if that call request succeeds and after the local NRM or ABM mode is entered.

When the SDLC PAD receives an I-Frame over the local SDLC interface, it transparently transfers the contents of the Information Field to the user data field of an unqualified Data packet which it transmits over the corresponding VC (with the correct LCN) to the remote end. If the size of the Information Field exceeds the maximum user data field for Data packets on that VC, the PAD transmits the Information Field in a complete packet sequence, conforming to the M bit procedures of CCITT Recommendation X.25 and section 4.3.4 of Part II of the IBM document cited earlier^[9].

2. Mapping is performed by the PAD between the SNRM(E), SABM(E), DISC, XID, TEST, UA, RD, DM, and FRMR frames on the local SDLC interface and corresponding QLLC qualified Data packets on the appropriate VC. I-Frames are handled as specified in the above item. All other SDLC frames are handled locally with no mapping to packets on the corresponding VC. With the exception of packets containing user data as specified in item 1, all other packet types (such as RNR and Reset Request) are handled (and possibly discarded) within the virtual circuit service with no

mapping to frames on the corresponding local SDLC interface. Tables 6-1 and 6-2 specify the relationships between the mappable frames and packets and summarize the actions taken by the PAD upon receipt of a specific frame or QLLC packet type. Table 6-1 summarizes frame formats and actions to be taken by the PAD upon receipt of each frame. Table 6-2 does the same for each QLLC packet type.

6.1.2 Differences if Type B Connection Supported Without QLLC

If the Type B connection support option applies to the VC, rather than QLLC (see section 3.2.2.1), section 6.1.1 applies, with the following differences:

1. With the exception of the mapping between unqualified Data packets and user data in call establishment/clearing packets on the VC and I-Frames on the local SDLC interface, none of the mapping specified in section 6.1.1 applies.
2. QLLC qualified Data packets should not be passed over the VC. Functions associated with these packets are handled locally, between the PAD and its directly connected SDLC DTE, as specified in section 3.2.2.1.

6.2 Send Queues and Flow Control: SDLC and Virtual Circuit

1. As discussed in section 3.2.2.3, the PAD should maintain a frame send queue for each active secondary station on (or emulated for) the local SDLC interface and a packet send queue for each virtual circuit into the network. An addition should be made to the appropriate secondary station SDLC frame send queue when a complete, valid, in-sequence frame is mapped from an incoming packet or packet sequence on the corresponding virtual circuit. When a complete, valid, in-sequence frame is received across the local SDLC interface, the PAD should be mapped it to one (or more, if segmenting is required) packets, which are added to the corresponding virtual circuit queue.
2. Virtual circuit send queues should be emptied in accordance with normal virtual circuit flow control procedures at the link and packet layers. For PAD SDLC interfaces at the primary station end, secondary station send queues should be emptied in response to polls of the corresponding station address from the primary. The entire queue should be emptied, with the last I frame or a following RNR frame having the F bit set. (If the queue is empty, the poll is acknowledged with an RR or RNR frame with the F bit set.) For PAD interfaces at the secondary station end, the send queues should be emptied in accordance with the poll procedures specified in section 3.2.2.3. The last frame should have the P bit set unless polling/SDLC procedures or flow control conditions preclude a response from the addressed secondary station. It is desirable that SDLC PADs be able to update the SDLC frame send queues to reflect newly added frames up until the time at which transmission of the frame currently last on the queue is started.
3. The PAD should handle sequencing and acknowledgment of the Data packets on the VC and of I-Frames on the local SDLC interface, independently, each in accordance with the procedures of the virtual circuit service and SDLC link interface, respectively. However, flow control conditions on the local SDLC interface and within the corresponding virtual circuit do, at least indirectly, influence one another. The PAD is responsible for mediating flow control between the two. For example, use of RNR and control of window advancement are available to the PAD to throttle input from the DTE on the local SDLC interface if network virtual circuit congestion occurs (control

of polling is also available on the interface at the secondary station end). If flow control for a secondary station on the SDLC interface becomes a limiting factor, the PAD can use packet layer RNR or window advancement strategies to throttle the corresponding virtual circuit.

6.3 PAD Actions in Response to Events & Received Frames/Packets

Section 6.1 specifies the various mappings between SDLC frames and virtual circuit packets that are to be performed by the PAD. This section specifies the other actions to be taken by the PAD in response to various events and the receipt of specific SDLC frames and QLLC packets. The first set of actions consist of clearing/resetting/re-establishing virtual circuits, deactivating/initializing SDLC stations, and/or transmitting certain frames over the local SDLC interface.

These specifications are presented by SDLC frame type and QLLC packet type. For each, the conditions under which the PAD transmits the frame or packet and the action taken by the PAD when it receives the frame or packet are specified. Distinctions are made, when appropriate, between PADs at the primary and secondary station ends, and between the QLLC and local SDLC support option cases.

1. Table 6-3 summarizes the conditions under which the PAD should transmit specific SDLC frames and the action(s) that should be taken by the PAD in response to the receipt of specific frames. Table 6-4 does the same for QLLC qualified Data packets. The handling of unqualified Data packets and packets other than Data packets has already been addressed in section 6.1.

A second set of special PAD actions is associated with the need to alert SDLC DTEs of changes in virtual circuit status. When a virtual circuit connection replaces a dedicated facility SDLC connection, it is important that the local polling function of the SDLC PAD(s) not prevent the DTE from receiving an indication of deleted PVCs, protocol problems, failed equipment, or other incidents of SVC clearing at the remote end. When there is an equipment failure or serious protocol violation at one access interface, the SDLC PAD at the remote end should receive some indication of the problem and take action at its SDLC interface to alert the DTE.

2. The SDLC PAD should modify its local polling and frame handling behavior for the affected SDLC secondary station address(es) during any of the following periods:
 - following deletion of a PVC, until a new VC is established and associated with the secondary station address
 - following resetting of a PVC with cause "out of order" or "network out of order," until subsequent resetting of the same PVC with cause "remote DTE operational" or "network operational," respectively
 - following clearing of an SVC, until a specified period of time has elapsed

A restart is treated as the associated set of resets and clears for the purposes of this discussion. At the secondary station end, the station address applies to the actual local secondary station locally attached. At the primary station end, the station address applies to the virtual secondary station being emulated for the remote end of the virtual circuit connection.

3. At the beginning of the period of modified operation, the PAD should request across the local SDLC interface that the affected secondary station be placed in disconnected mode, if it is not already in that mode. At the secondary station end, the PAD should

transmit a DISC and at the primary station end, the PAD should transmit a RD when permitted and respond to a DISC command. Following this initial action, and until a new VC is established and associated with the affected secondary station (or the established PVC is again declared operational), the PAD should not transmit any frames with the affected secondary station address across the local SDLC interface, unless the frame is part of the process of declaring the VC operational again. Following the period of modified operation, the PAD should return to normal polling and frame handling operation and attempt to leave disconnected mode, if necessary.

The precise action the PAD takes during the period of modified operation, after the initial disconnection mode request, depends on the type of VC disruption that has occurred and whether the local or remote end initiated the VC disruption.

If the VC disruption was caused by the deletion of a PVC or by the PVC being declared not operational from the remote end, the PAD should not transmit any frame with the associated station address across the local SDLC interface until a new VC is provisioned/established or the PVC is declared operational again for that station address. If the VC disruption was caused by the PVC being declared not operational by the local end PAD, that PAD may transmit and respond to frames across the local interface necessary to declare the PVC operational again (e.g., if a lack of response to PAD polls at the secondary end resulted in the non-operational declaration, the PAD may attempt to poll the station with SNRMs at a reduced rate to detect when the station returns to operational status).

If the VC disruption was caused by the clearing of an SVC, all frame handling associated with the affected SDLC station should be halted by the SDLC PAD for the period T_{d1} if the clearing was initiated by the network or the remote end and for the period T_{d2} if the clearing was initiated (or specified) by the local DTE. During this period, all attempts to establish new virtual circuit connections to or from the affected secondary station are prevented. Each of these two values is an interface configuration option which should be capable of being set to at least the values in the range of 0-120 seconds. Following this buffer period, which alerts the local DTE of the clearing at the SDLC protocol level, normal polling, frame handling and virtual circuit establishment procedures are resumed by the SDLC PAD.

6.4 Exception Conditions, State Changes, & Diagnostics

The PAD should monitor the status of the physical level and each link level station for each local SDLC interface and appropriately reflect state changes upon receipt of relevant local SDLC interface frames and/or packets on the corresponding virtual circuit. The PAD should also take appropriate action in response to state or status changes. When virtual circuits are cleared or reset by the PAD, the appropriate cause and diagnostic codes should be used. The following specific requirements apply:

1. The PAD should change the status of a local SDLC station to (normal/asynchronous) disconnected mode (dm) when it receives/sends a DM, FRMR, SNRM(E), or SABM(E) frame or it receives/sends a UA frame response to a DISC with the corresponding station address across the local interface.
2. The PAD should change the status of a local SDLC station to normal response or asynchronous balanced mode, as appropriate, when it sends or receives a UA frame with the corresponding station address across the local interface, in response to a previous SNRM(E) or SABM(E).

3. The PAD should clear an SVC (reset a PVC) and purge all as yet undelivered packets and deliver (if possible) or purge frames associated with that VC in response to any of the following:
 - A network problem or other condition occurs at the packet layer that requires clearing or resetting according to Recommendation X.25 (e.g., invalid packet type or packet too long).
 - The conditions specified in item 2 of section 5.4.
4. If a reset signals a transition between DTE or network "operational" and "out of order" states, the PAD should reflect the change across the local SDLC interface (see SM, DISC, DM, and RD actions taken in response to "oper." and "nonoper." in Table 5-2).
5. See section 6.3 for the PAD actions to be taken in response to the clearing of an SVC or the "out of order" resetting of a PVC.
6. The PAD at the secondary station end should attempt to restore the station to normal response or asynchronous balanced mode (NRM/ABM) if it is in the disconnected mode, unless the station or interface has been declared unoperational (because the limit has been reached on attempts to restore the station, or for another reason) or the period of modified operation defined in items 2 and 3 of section 6.3 applies.
7. For PVCs, the PAD should initiate and respond to resets as specified in section 6.3 above and in section 4.4.3 of Part II of the IBM document including the QLLC specifications^[9]. Resets should cause the PAD to issue a DISC or RD to the corresponding SDLC station across the local interface in all cases except for the causes "network operational", "remote DTE operational", or "DTE originated" (with diagnostic decimal 0). In these three cases, the PAD on the secondary end should attempt to restore the associated station to the NRM/ABM if it is not already in that mode.
8. When a PAD issues a Clear Request, Reset Request, or Restart Request on a connection supporting QLLC, it should use the diagnostic codes specified in Appendices E and F of the IBM QLLC document cited earlier^[9]. When the PAD is acting on behalf of the SDLC DTE (if a native X.25 DTE would have initiated the action on the virtual circuit if it were present in place of the PAD and attached SDLC DTE), then the DTE diagnostics of Appendix F of that IBM document should be used.
9. The procedure for attempting to re-establish an SVC that is specified in item 9 of section 5.4.2.1 should be used by the PAD following those cases of clearing specified in that item.

Table 6-1. QLLC Mapping: From Local SDLC Interface Frames

Frame Label	QLLC Qualified Data Packet Sent Across VC by PAD, Upon Receipt of SDLC Frame	
	At Primary Sta. End	At Secondary Sta. End
SNRM(E) or SABM(E)	QSM	NA
DISC	QDISC	NA
XID	QXID (C)	QXID (R)
TEST	QTEST (C)	QTEST (R)
RD	NA	QRD
UA	NA	QUA
DM	NA	QDM
FRMR	NA	QFRMR

Legend & Notes

- (C) Command version (originated by primary station end)
- (R) Response version (originated by secondary station end)
- NA Not Applicable or Invalid

Notes:

1. Address field of each frame applies to secondary station (may be broadcast x'FF' value); each active secondary station uniquely maps to a QLLC virtual circuit LCN.
2. Frame control field and information field (if present) conform to SDLC specifications; information field maps to/from corresponding QLLC packet fields.

Table 6-2. QLLC Mapping: From Virtual Circuit Qualified Data Packets

QLLC Qualified Data Packet				Frame Sent Across Local SDLC Interface by PAD, Upon Receipt of QLLC Packet	
Label	Address Field*	Control Field (Hex)	Info. Field**	At Secondary Sta. End	At Primary Sta. End
QSM	FF	93	No	SNRM(E) or SABM(E) %	NA
QDISC	FF	53	No	DISC	NA
QXID (C)	FF	BF	Yes	XID	NA
QXID (R)	xx	BF	Yes	NA	XID
QTEST (C)	FF	F3	Yes	TEST	NA
QTEST (R)	xx	F3	Yes	NA	TEST
QRR (C)	FF	F1	No	RR	NA
QRR (R)	xx	F1	No	NA	RR
QRD	xx	53	No	NA	RD
QUA	xx	73	No	NA	UA
QDM	xx	1F	No	NA	DM
QFRMR	xx	97	Yes	NA	FRMR

Legend & Notes

- (C) Command version (originated by primary station end)
- (R) Response version (originated by secondary station end)
- NA Not Applicable or Invalid

Notes:

- * FF is broadcast address used in commands; xx is secondary station address value other than FF.
- ** If present ("Yes"), copied to/from corresponding SDLC frame.
- % Use of SNRM, SNRME, SABM, or SABME depends on SDLC interface configuration.

Table 6-3. SDLC Frames: Summary of PAD Send Conditions & Actions on Receipt

Frame Type	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
SNRM(E) or SABM(E)	Never	When (re-)initializing** local link; mapped incoming QSM packet if QLLC supported	If QLLC not supported: respond with UA (DM if PAD interface not operational*); station set to dm. If QLLC supported: map to QSM.	Discard, Re-Initialize SDLC Station or DISC if beyond retry limit, Clear(Reset) SVC(PVC)
DISC	Never	In response to RD; to disconnect secondary station; mapped incoming QDISC if QLLC supported	If QLLC not supported: respond with UA; station set to dm. If QLLC supported: map to QDISC.	Discard, Re-Initialize SDLC Station or DISC if beyond retry limit, Clear(Reset) SVC(PVC)
XID	See section 5.3	See section 5.3	See section 5.3	See section 5.3
TEST	As response to TEST command if QLLC not supported; mapped incoming QTEST packet if QLLC is	Mapped incoming QTEST packet if QLLC supported	Send TEST response if QLLC not supported; Map to QTEST if QLLC supported	Accept & discard if QLLC not supported; Map to QTEST if QLLC supported
RR	Window update poll response if not done by I-frame and no receive flow control	Window update (poll) if not done by I-frame and no receive flow control	Update window; end transmit flow control state if appropriate	Update window; end transmit flow control state if appropriate
RNR	To signal receive flow control	To signal receive flow control	Update window; flow control send	Update window; flow control send

* DM is used as response at primary end whenever not operational or remains in dm.

** As used here, "initialize" means to put secondary station in dm and then attempt to restore to NRM or ABM.

Note: Above does not include responses required solely by P/F bit rules.

Table 6-3. SDLC Frames: Summary of PAD Send Conditions & Actions on Receipt (Cont.)

Frame Type	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
REJ	Never	Never	Update window, (re)transmit I frames as requested	Update window, (re)transmit I frames as requested
RD	If PAD interface is inoperative; non-normal SVC clear; mapped incoming QRD if QLLC supported	Never	Respond with FRMR	If QLLC not supported: respond with DISC. If QLLC supported: map to outgoing QRD.
UA	In response to all commands required by SDLC; mapped incoming QUA if QLLC supported	Never	Respond with FRMR	Update station status, if in response to SNRM(E), SABM(E) or DISC; map to QUA if UA in response to command initiated by incoming QLLC packet. If QLLC not supported, take action on VC to reflect new status, if appropriate (e.g., Table 5-2).

** As used here, "initialize" means to put secondary station in dm and then attempt to restore to NRM or ABM.

Note: Above does not include responses required solely by P/F bit rules.

Table 6-3. SDLC Frames: Summary of PAD Send Conditions & Actions on Receipt (Cont.)

Frame Type	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
DM	If not operational or no VC established, or when local station in dm*, in response to non-mode-setting command; mapped incoming QDM if QLLC supported	Never	Respond with FRMR	Note station mode (do not re-initialize** if attempt limit exceeded); map to QDM if DM in response to command initiated by incoming QLLC packet
FRMR	In response to invalid/unsupported frame when in NRM/ABM mode; mapped incoming QFRMR if QLLC supported	Never	Respond with FRMR	Re-Initialize** SDLC Station, Clear(Reset) SVC(PVC)
I	Mapped unqualified Data packet or call setup/clear packets with user data (see sect. 6.1.1) in response to poll	Mapped unqualified Data packet or call setup/clear packets with user data (see sect. 6.1.1) according to poll procedures	Map to unqualified Data packet over VC	Map to unqualified Data packet over VC
Other	Never	Never	Respond with FRMR	Discard, Re-Initialize SDLC Station, Clear(Reset) SVC(PVC)

* DM is used as response at primary end whenever not operational or remains in dm.

** As used here, "initialize" means to put secondary station in dm and then attempt to restore to NRM or ABM.

Note: Above does not include responses required solely by P/F bit rules.

Table 6-4. QLLC Packets*: Summary of PAD Send Conditions & Actions on Receipt

QLLC Packet	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
QSM	Mapped SNRM(E) or SABM(E)	Never	Map to SNRM(E) or SABM(E) on local interface (let primary station detect this error by remote end)	Map to SNRM(E) or SABM(E) on local interface; station set to dm
QDISC	Mapped incoming DISC	Never	Map to DISC on local interface (let primary station detect this error by remote end)	Map to DISC on local interface and update station status to dm
QXID	See section 5.3	See section 5.3	See section 5.3	See section 5.3
QTEST	Mapped incoming TEST	Mapped incoming TEST	Map to TEST on local interface	Map to TEST on local interface
QRR	Never	Never	Update window and flow control receive accordingly	Update window and flow control receive accordingly
QRD	Never	Mapped incoming RD	Map to RD on local interface	Map to RD on local interface (let secondary station detect error and respond with FRMR)

* Assumes QLLC supported. If QLLC not supported: (a) none of these packets is transmitted except for UDP and (b) if any of these packets (aside from UDP) is received, the response is a QFRMR from the secondary end or a QDISC followed by clearing of the VC from the primary end.

Table 6-4. QLLC Packets*: Summary of PAD Send Conditions & Actions on Receipt (Cont.)

QLLC Packet	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
QUA	Never	Mapped incoming UA	Map to UA on local interface and update station mode based on frame responding to; take action on VC to reflect new station status (e.g., Table 5-2)	Map to UA on local interface (let secondary station detect error and respond with FRMR)
QDM	Never	Mapped incoming DM	Map to DM on local interface and update station status	Map to DM on local interface (let secondary station detect error and respond with FRMR)
QFRMR	Never	Mapped incoming FRMR	Map to FRMR on local interface, clear(reset) SVC(PVC), and update link status	Map to FRMR on local interface (let secondary station detect error and respond with FRMR)
UDP (Data)	Mapped incoming I-Frame; update window	Mapped incoming I-Frame; update window	Map to I-Frame on local interface; update window	Map to I-Frame on local interface; update window

* Assumes QLLC supported. If QLLC not supported: (a) none of these packets is transmitted except for UDP and (b) if any of these packets (aside from UDP) is received, the response is a QFRMR from the secondary end or a QDISC followed by clearing of the VC from the primary end.

Appendix A

Description of the SNA Upper Layers Processor Concept

The proposed generic requirements for PPSN support of SNA/SDLC equipment specify network PADs that protocol convert between an X.25 type virtual circuit connection within the network and an SDLC access interface. The higher layers of the SNA protocol stack are passed transparently as user data by the PAD. This approach permits network transparency for such functions as customer network management and a relatively cost-effective PAD implementation. However, certain services cannot be provided within the PPSN if only the SDLC layer is processed:

- supporting separate virtual circuit connections (i.e., simultaneous access to different host computers) for individual terminals on a single cluster controller
- providing asynchronous terminal to 3270 Data Stream protocol conversion

If provision of one or more of these services by the PPSN is warranted, the question remains of the most effective approach for rendering such services. Providing such higher layer SNA functionality in each PAD may be inadvisable because:

- demand for these services may vary widely from PAD to PAD
- implementation of these services will likely represent a heavy processing burden for PAD type devices
- software updates to accommodate additional services and/or IBM SNA modifications would be difficult to administer across a large number of distributed PADs of multiple vendors

For these reasons, the concept of centralizing such higher layer SNA functions and services in a limited number of specialized processors within the PPSN is attractive. In addition, separation of this function from the PAD permits the basic support of SNA equipment by the PPSN to be specified without the additional delay that would result if the PAD were required to implement all of these functions.

The remainder of this appendix presents a high-level description of the concept of PPSN processors or nodes designed to implement services associated with layers of the SNA protocol stack above the SDLC link layer. The description focuses on the service of switching different terminal devices (LUs) on a single cluster controller (PU) to applications on different hosts. It is intended to clarify what is meant by separating such functions from the PAD implementations and to prompt technical comments and suggestions from interested parties.

Figure A-1 illustrates two cluster controllers, each with two terminal devices and each connected to the PPSN via an SDLC PAD. For each cluster controller, one terminal wishes to access an application on a host II (via a NPSI X.25 connection to the PPSN) and the other terminal wishes to simultaneously access an application on a host III (via an SDLC connection to a PAD on a remote PPSN). Since this splitting of LUs on the same PU requires switching based on layers of the SNA protocol stack above the link layer, the services of a PPSN SNA Services Processor (PSSP) are proposed.

Instead of establishing a virtual circuit connection directly between each cluster controller and a desired host FEP, a two-stage procedure is used. The virtual connections could be either SVCs or PVCs. For illustrative purposes, the following assumes all connections are SVCs. Assuming the virtual connection is initiated from the terminal end, an SVC is established from

each of the two cluster controllers to the PSSP (i.e., the called address in the Call Request packet is that of the PSSP, which would likely be part of an X.25 hunt group). The PSSP is the destination specified whenever its services are required. Once this virtual circuit connection between the customer cluster controller and PSSP is established, the PSSP intercepts information carried at higher SNA layers to switch individual LUs to different hosts, over SVCs established between the PSSP and each of the host FEPs involved. The processor would also maintain records necessary to bill the customer for all virtual circuits established and the SNA services provided.

Figure A-1 suggests the outlines of a possible approach for performing this LU switching function. Toward the secondary station (terminal) end, the PSSP acts as an SSCP (I). Toward the primary station (host) end, the PSSP acts as a collection of virtual (emulated) secondary PUs and associated virtual LUs. The PSSP associates each LU on the terminal end with a virtual LU belonging to one of the virtual PUs being emulated toward the host end. On both ends, there is a one-to-one correspondence between PUs and virtual circuits. The number of active LUs at the terminal end and active virtual LUs being emulated toward the host end are the same. However, the number of virtual PUs emulated toward the host end and the grouping of LUs per PU in general differ from that at the terminal end.

The PSSP performs the switching function by mapping between actual PU/LU combinations at the terminal end and virtual PU/LU combinations emulated toward the host end. This mapping is performed on the basis of: (a) information stored in the PSSP for each calling DTE desiring this service and (b) LU name and SNA NAU address information carried in SNA Path Information Units (PIUs). The calling DTE address in the Call Request packet and the short form of the originating address field (OAF') in a PIU identifies the terminal end LU. Using this information and the "name of the other LU" in the INIT-SELF request RU, the PSSP would be able to look up (in its configuration tables) the DTE called address of the virtual circuit connection and the short form destination address field (DAF') identifying the owning SSCP for the host end.

A subsequently received BIND would contain the OAF' identifying the address that would identify the host end (application) LU to be used in LU-LU session RUs for the terminal end LU identified, until this session is brought down. Once a session (e.g., SSCP-LU or LU-LU) was established, the virtual circuit logical channel number (LCN) and OAF'/DAF' values in the PIU at one end of the PSSP would determine the LCN and OAF'/DAF' values to use in mapping to the opposite end.

The above discussion is intended to introduce the concept of a network processor for higher layer SNA services and sketch some possibilities for its operation. It is intended to alert the reader to possibilities under consideration for future generic requirements. Comments and suggestions are encouraged.

LU 1 TALKS TO LU y APPL. IN HOST III
 LU 2 TALKS TO LU w APPL. IN HOST II

○ = LU
 □ = PU

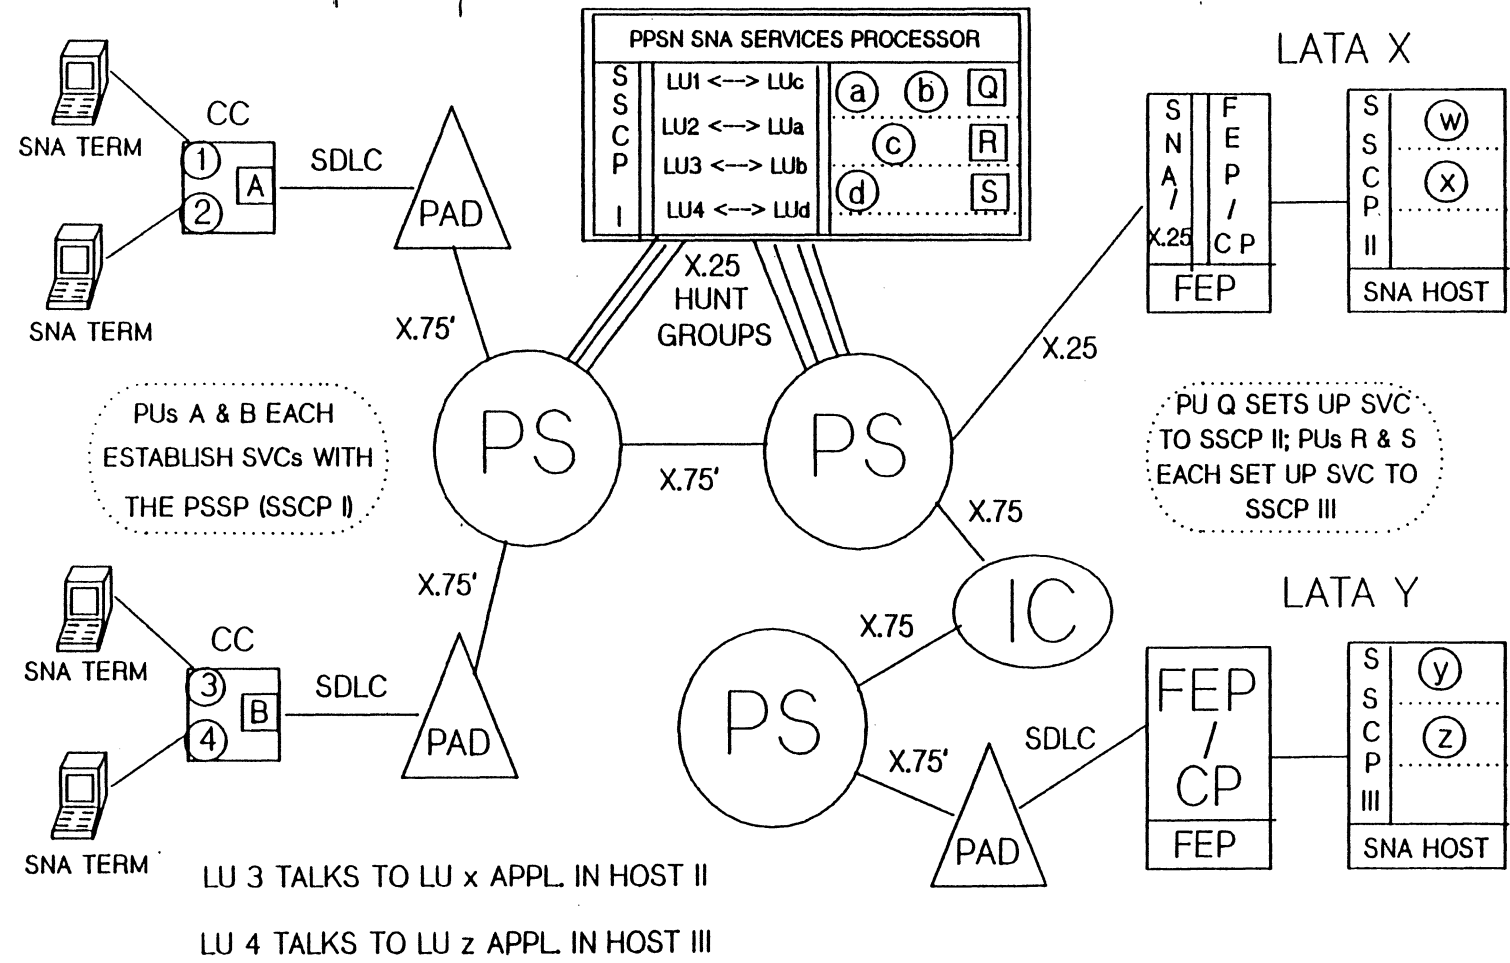


Figure A-1. Illustration of SNA Services Processor Concept for LU Switching

BLANK PAGE

Appendix B

Prototype Cases Used for Performance Specifications*

Assumptions and Defined Quantities

Assume:

1. Secondary station releases control after full block or transfer
2. For file transfer, all data is from secondary to primary
3. For interactive sessions, each secondary to primary transfer followed by primary to secondary transfer of same size
4. Packet handling capacity over VCs supporting the local SDLC interfaces is not a limiting factor

N_1 = Assumed no. of SDLC interfaces active on AC or PS (should be > 15)
= Maximum no. that can be configured

N_{smax} = Maximum no. of secondary stations per SDLC interface that can be supported (must be at least values specified in sect. 3.2.1, item 4)

$N_{sconf} = 0.9 * N_{smax}$ = Ave. no. of secondary stations configured per SDLC intf.

$N_{abusy} = INT(0.75 * N_{smax})$ = Assumed ave. no. of active secondary sta. per SDLC intf. (busy hour)

$N_{ibusy} = INT(0.15 * N_{smax})$ = Assumed ave. no. inactive secondary sta. per SDLC intf. (busy hour)

$N_{amod} = INT(0.3 * N_{smax})$ = Assumed ave. no. of active secondary sta. per SDLC intf. (mod. traffic period)

$N_{imod} = INT(0.6 * N_{smax})$ = Assumed ave no. of inactive secondary sta. per SDLC interfaces (mod. traffic period)

$I_{max} = 265$ octets = Assumed max. size of SDLC I frame information field

$M_{file} = 25000$ octets = Assumed size of average file transfer

$M_{big} = 2000$ octets = Assumed size of large interactive block

$M_{ave} = 50$ octets = Assumed size of average interactive transfer

T_{max} = Max. throughput in bps AC or PS can handle across all configured interfaces

$T_{high} = 0.40 * T_{max} / N_1 / N_{sconf}$ = Assumed per-sta. average bps (high)

$T_{mod} = 0.15 * T_{max} / N_1 / N_{sconf}$ = Assumed per-sta. average bps (moderate)

S_{max} = Max. access line speed supported on AC or PS (9.6, 19.2, or 56 kbps)

* See item 2 of Section 3.2.2.3 for relevant performance specification.

(Continued)

Case 1: Single Very Large File Transfer User; Other Stations Low Activity

I_{max} applies here and in all cases in this table

N_{ibusy} stations per line dormant

$N_{abusy} - 1$ stations per line responding to every fifth poll (evenly distributed over cycles)

Data transfer of M_{ave} octets in both directions with each active cycle

Line speed = 1200 bps

One station per line in file transfer mode; secondary to primary direction:

Responds with one block for each poll cycle, sent at max. rate

Sends block of M_{file} octets and then releases control

Line speed = S_{max}

Case 2: All Active Stations Operating at Low Level Consistently

N_{ibusy} stations per line dormant

N_{abusy} stations per line respond to each poll for which transfer block ready:

M_{ave} octets per transfer

Messages accumulate at secondary at rate of T_{mod} bps

Primary responds with message of equal size

Case 3: Moderate Number of Stations Active; Imbalance in Activity Levels

N_{imod} stations per line dormant

One secondary station per line in file transfer mode:

When message block complete, responds to poll and then releases control

M_{file} octets per block

Accumulates block at rate of T_{high} bps

Line speed = S_{max}

$INT(N_{amod}/2)$ stations per line operating at high interactive level:

Responds to each poll for which transfer block ready

M_{big} octets per transfer

Messages accumulate at rate of T_{high} bps

Primary responds with message of equal size

Line speed = 9600 bps

Remaining active stations operating at low interactive level:

— Responds to each poll for which transfer block ready

M_{ave} octets per transfer

Messages accumulate at rate of T_{mod} bps

Primary responds with message of equal size

Line speed = 1200 bps

Appendix C

Asynchronous Control Interface Protocol Extension

Overview and Conventions

Defined below for those asynchronous interfaces configured for control of SDLC PAD SVCs are:

- a. an escape mechanism for toggling between the normal (X.28/X.29) and SDLC PAD control function modes for the asynchronous interface; and
- b. details of the command format used when the SDLC PAD control function mode is active.

When the SDLC PAD control function mode is NOT active, the asynchronous protocol and command/display formats specified in section 3.3 of the PPSNGR apply. The formats used by network control interfaces to prompt the user, acknowledge command requests, and display SVC status information are specified in Appendix D.

In this Appendix and in Appendix D, the following conventions apply:

- described (non-literal) entries or fields are enclosed in angle brackets ("`< description >`")
- options are enclosed in square brackets ("`[options]`")
- commas separate options within square brackets which are not mutually exclusive ("`[A, B, C]`") and the commas are included to separate any options actually selected in the command itself
- pipes are used to separate mutually exclusive alternatives ("`A|B`")
- all other characters and special symbols (including parenthesis) are literal and are entered (or displayed) exactly as illustrated

Escape Sequence Mechanism to Enter/Leave SDLC PAD Control Function Mode

The exclamation point character ("!") is used to toggle between normal asynchronous and SDLC PAD control function modes. The toggling action does not become effective until a carriage return ("`<CR>`") is entered. When in the control function mode, the system prompts the user with the string "`_SDLCntl>`" as a mode indicator. The system also precedes any SDLC PAD SVC message with the same string. (see Appendix D). When in the control function mode, SDLC PAD SVC requests can be entered, using the formats specified later in this table.

A request may be entered on the same line as the escape character, as illustrated below (assuming we are initially in the normal asynchronous mode):

```
!<SDLC PAD SVC request>[!]<CR>
```

The "!" as the first character indicates escape to the command mode. The entire command is terminated by a carriage return. Once in the SDLC PAD command mode, all subsequent entries are assumed to be in this mode until the escape is terminated with a line ending in a "!" plus carriage return. If only a single control mode request is needed, escape into and out of the control mode can be accomplished in a single line, as illustrated.

Format of SDLC PAD SVC Establishment and Secure Dial-Back Requests

In all cases except for a dial-in asynchronous interface with the secure dial-back feature, the format of a call request is:

<Field unique to SDLC PAD SVC>[,<SVC facil.>]-<dest. address>[*][D|P<user data>]<CR>

The request format consists of two major parts: (a) a field that is unique to the needs of a SDLC PAD SVC and (b) the optional facility, destination address, and user data fields associated with a normal asynchronous call request (see PPSNGR section 3.3). This approach makes maximum use of the existing asynchronous interface protocol and formats, while introducing those new elements that are necessary for SDLC support and establishing a call that originates on another interface.

In the case of a dial-in interface, the system dials back to the address established for the authorized user for security purposes. The format used on the initial dial-in to request dial-back is:

Z<auth. code><CR>

The format of the field unique to the SDLC PAD SVC in the call request is:

Z<auth code>[,#<VC prof.>][,I<start>][,E<end time>][,A<orig. addr.>][,X<hex CUD>][,Y<trans. del.>]

There is one required and several optional subfields unique to the SDLC PAD SVC request, defined as follows:

1. Each SDLC PAD SVC request starts with an authorization code, preceded by the code letter "Z". This same code is used in requesting secure dial-back. The authorization code, which is separate from any NUI selection applying to the SVC, validates that the requester is authorized to set up an SVC on behalf of one or more (other) interfaces. For non-switched interfaces, the authorization code is an alphanumeric string of 6 to 10 characters which must be unique within the network element. For switched interfaces, the code should be 10 characters and unique within the PPSN. In both cases, the code must contain at least one decimal digit or special symbol and at least one letter, and must not have more than three instances of the same character. (The network implementation of the control interface turns off local echo between the "Z" code letter and the <CR>, comma, or dash which terminates the code value.)

Each non-switched asynchronous control interface should be configurable with at least two valid authorization code values and each switched control interface should be configurable with at least 100 valid authorization code values. Associated with each authorization code that is configured for a control interface is a list of DTE addresses which the control interface user associated with that code may specify as the calling address for the SVC. Addresses of grouped PU/LU characteristics interfaces and addresses which serve only as a hunt group identifier should not be used as such a calling address.

This list should be able to accommodate at least five full addresses or truncated X.121 addresses. A truncated X.121 address (no more than 2 least significant digits missing) is assumed to match any full address with the same first digits. [Such truncated addresses would be useful when a customer wishes a single control interface to be able to control all interfaces operated by customer employees within a tens or hundreds number group block.] Addresses (with or without escape codes) will conform to section 4 of the PPSNGR. Also associated with each authorization code for dial-in control interfaces will be an address that the network can use for secure dial-back before any control mode request is honored.

2. Optionally, a virtual circuit profile number may be specified following the "#" code symbol (see section 5.5 for description and format). If present, this profile, stored by the network, provides specifications for the SVC to be established, including values for both regular X.25 Call Requests and those added details unique to SDLC PAD SVCs. Any optional entry explicitly present in the SVC request adds to or overrides the corresponding value of the profile.
3. Optionally, a starting time for the SVC may be specified following the code letter "I". If none is specified, the call is initiated as soon as possible. Optionally, a time for clearing the SVC may be specified following the code letter "E". If none is specified the call is cleared only by explicit subsequent command or a failure condition. All time values are specified in the following format:
$$[&|_]hh:mm:ss$$
where "hh" is a 24-hour hour value, "mm" is minutes, and "ss" is seconds. The time value is an absolute time value during the next 24 hours, unless preceded by an optional "&" or "_" (underscore) character. A "&" indicates that the value is relative to the current time being 00:00:00. A "_" indicates that a call is to be cleared after the designated duration has passed with no unqualified Data packets being passed over the SVC in either direction.
4. Optionally, the address of the DTE from which the Call Request is to be initiated is specified following the code letter "A". This address value has the same format as the called address field in the normal (X.28) asynchronous call setup command (see section PPSNGR section 3.3). If no calling address is provided, the first non-truncated address appearing in the calling address list for the authorization code is assumed.
5. Each type A connection SVC establishment request (either explicitly, or via a VC profile designation) includes a specification of the first 1 to 15 octets of the call user data (CUD) field, expressed as twice that number of hexadecimal digits (semi-octets), following the code letter "X". These octets precede any user data specified in the call user data field at the end of the normal portion of the X.28 call setup command. These initial CUD octets are used to specify the protocol identifier, and secondary station PU/LU characteristics, as specified in item 3 of section 5.4.2.1.
6. Optionally, the desired maximum transit delay for the SVC is specified, in milliseconds, following the code letter "Y" (1-65534).

The following illustrates a SDLC PAD SVC establishment request, assuming we are already in the command mode:

```
SDLCntl> Z5551234JM$,E_00:05:00,A5551111,XC3010000,T3299,R-132505559999*Ppasswd<CR>
```

Above, a user with authorization code 5551234JM\$ (not echoed when entered) requests that an SVC be established from calling party 555-1111 to called party 555-9999 on destination network 3250, via the IC with DNIC 3299. The SVC is to be reverse charged, cleared within five minutes of no data traffic being detected, and the CUD field should contain four octets with the hex value "C3010000" followed by six ASCII-encoded octets, "passwd".

SDLC PAD SVC clear request format

While in the SDLC PAD control function mode, the user may request that any existing SVC for which it has control authorization be cleared. This clearing takes precedence over any pre-defined clearing time specified in the SVC establishment request. The format of the clear request is as follows:

```
CLR|clr:<address of calling DTE>,<logical channel number>[P|D<clr user data>]<CR>
```

where the originating DTE address and logical channel number uniquely identify the SVC, from among those that are currently active and that this user is currently authorized to control. The address and logical channel number are identical to those displayed by the network in its message confirming establishment of an SDLC PAD SVC (see Appendix D). Optionally, clear user data in octet-encoded ASCII format may also be specified. Clear Request packets generated as a result of such commands have the cause "DTE originated" and diagnostic code #0.

Message Display Request Formats

The default is to display status messages for any SVC established by the control interface as long as that interface is functioning. However, two commands are available to turn off message display on the interface (e.g., to avoid confusion while using the interface for a regular asynchronous session) and to re-enable messages:

```
MSG|msg on|off<CR>
```

Appendix D

Control Interface Display Formats for Status and Clearing Messages

The control interface which establishes an SDLC PAD SVC will receive all status and clearing messages associated with that SVC for as long as that control interface is maintained, unless message display has been disabled (using the "MSG off" command). SDLC PAD control messages will be distinguished from those sent to normal (non-control) asynchronous terminals by the string "SDLCntl> " that precedes all messages relating to SDLC PAD SVCs and serves as the system prompt when the interface is in the control function mode. When the interface is not in the control mode, the normal asynchronous prompt (e.g., "**") is used.

Any message applying to an established SVC will be preceded by

<calling DTE address>,<logical channel number>:

as two decimal numbers, separated by a comma and terminated with a colon (for example, "SDLCntl> 2015551212,003: REQUESTED SVC ESTABLISHED"). The following are the messages to be used in communicating with an asynchronous control terminal:

Messages Not Applying to an Established SVC

INVALID COMMAND SYNTAX (<invalid command string>)

INVALID FIELD VALUE OR VALUE COMBINATION (<offending fields>)

AUTHORIZATION CODE INVALID OR MISSING

CODE NOT AUTHORIZED FOR SPECIFIED CALLING PARTY

SPECIFIED CALLING DTE NOT ACCESSIBLE

PREPARE FOR SECURE DIAL-BACK; DIAL-IN CONNECTION BEING BROKEN

REQUESTED DIAL-BACK; PLEASE ENTER SDLC CONTROL COMMAND(S):

CALLING END BUSY; CALL REQUEST FAILED

CALLED END BUSY; CALL REQUEST FAILED

CALL REQUEST FAILED (<X.25 cause>, <X.25 diagnostic decimal code>)

CLEAR REQUEST FAILED

SPECIFIED SVC DOES NOT EXIST

Messages Applying to an Established SVC (Preceded by Address,LCN)

REQUESTED SVC ESTABLISHED

SVC ESTABLISHED, BUT TO ALTERNATE CALLED ADDRESS (<new called address>)

REQUESTED SVC CLEARED

SVC CLEARED BY NETWORK (<X.25 cause>,<X.25 diagnostic decimal code>)

SVC CLEARED BY DTE (<X.25 cause>,<X.25 diagnostic decimal code>)

SVC RESET (<X.25 cause>,<X.25 diagnostic decimal code>)

NEW SVC CANNOT BE ESTABLISHED DURING BUFFER PERIOD (<time remaining>)

CALL RE-ESTABLISHMENT ATTEMPT UNDERWAY

Appendix E

Recommended Formats for Menu-Oriented Command & Display Screens

Conventions and General Rules

1. Fields on the screen that are to be provided by the user appear in italics. If no value is entered into a field, a default is used if one is specified. If a VC profile is specified and it contains a value for the parameter involved, it becomes the default. Otherwise, the default indicated in curly brackets ({}) is used, if one is specified.
2. Descriptions of entries appear in angle brackets (<>). Optional values appear in square brackets ([]). Mutually exclusive alternatives are separated by a pipe symbol (e.g., "A|B" = A or B). Editorial notes or explanations that do not actually appear on a screen are enclosed by slashes (e.g., /note/). Everything else is literal.
3. In moving around a screen and changing screens, the following rules apply:
 - a carriage return, enter, or tab is used to move to the next field
 - after the last data field, the cursor moves to the help/action field, if it is present
 - a back tab moves to a previous field
 - "cntl a" moves the cursor to the beginning of the field
 - "cntl e" moves the cursor to the end of the field
 - "cntl k" erases the current field
 - backspace ("cntl h") moves the cursor to the previous field character position
 - over-writing a character replaces it within the field
4. Within a field, entering "?" or "<ESC>?" displays the appropriate format for that field on the assistance message line. The latter alternative is used when a "?" would be a valid entry within the current field.
5. The sequence "<ESC>h" moves the cursor to the help/action field of the screen.
6. With the exception of the authorization screen, after the last blank field of a screen has been filled in and after the "P" command has been entered to process the screen, display of the same screen (with the same field values) is retained. If the screen is valid, the cursor is placed in the help/action field. If the screen has any detected errors, the cursor is placed at the beginning of the first field in which an error was detected. However, syntax errors should be diagnosed on the error/assistance message line of the screen at the time that field is first entered.

SDLC PAD CONTROL MODE
SCREEN #0: AUTHORIZATION

[New] Authorization Code: *<6-10 characters>*

<Entry Error Diagnostics & Assistance Messages Are Displayed Here>

- /Note 1: Word "NEW" does not appear for initial authorization/*
- /Note 2: After 3 invalid entries, forced exit from control mode or disconnect/*
- /Note 3: After valid entry, sent to screen #1 for non-switched or dial-back interface/*
- /Note 4: After valid entry, execute dial-back for dial-in interface/*
- /Note 5: Authorization code is not echoed on screen./*

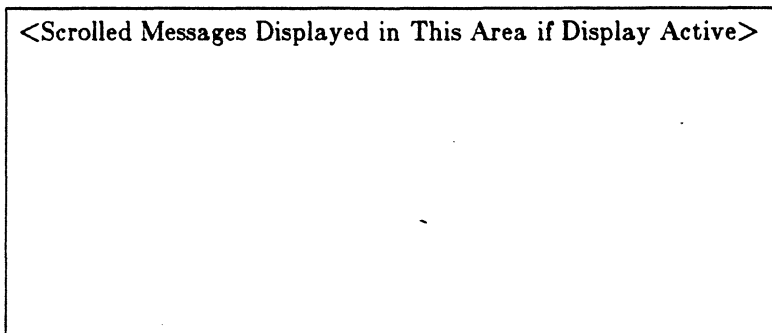
SDLC PAD CONTROL MODE

SCREEN #1: REQUEST/COMMAND SUMMARY

- 0: Authorization
 - 1: Request/Command Summary
 - 2: SVC Establishment
 - 3: SVC Clear
 - 4: <Reserved>
 - 5: <Reserved>
 - 6: <Reserved>
 - 7: <Reserved>
 - 8: <Reserved>
 - 9: Help/Format Summary
- E: Exit Control Mode

Screen Desired: <0-9 or E>

<Scrolled Messages Displayed in This Area if Display Active>



Help & Screen Action Requests: /0-9, C, E, F, P, N, R, or T/

<Entry Error Diagnostics & Assistance Messages Displayed Here>

/In help & action field (for all screens)

- 0-9: Go immediately to screen specified
- C = clear all entries on this screen
- E = exit control mode
- F = turn message display off
- P = process the screen as currently filled out
- N = turn message display on
- R = return to last entry field worked on above
- T = go to top of this screen (first field) /

SDLC PAD CONTROL MODE
SCREEN #2: SVC ESTABLISHMENT

VC Profile #: <1 or 2 digit value>

Start time: /@/<hh:mm:ss>{800:00:00} End time: /@_/<hh:mm:ss>

Orig. addr.: <PPSN addr. or abbrev.> Dest. addr.: <X.121 addr. or PPSN addr./abbr.>

Desired Carrier (RPOA) DNIC(s): <0 to 4 DNIC(s), separated by commas>

NUI Billing/Profile Customization ID/PIN: <first field>/,<second field>/

CUG Index: <1, 2, or 4 digit value> Outgoing access?: Y|{N}

Reverse Charge?: Y|{N} Charging Information?: Y|{N}


Packet size: O->D: {128}|256,512 D->O: {128}|256,512

Packet window size: O->D: <1-7 or 1-127>{2} D->O: <1-7 or 1-127>{2}

Throughput class: O->D: <75,...,48000>{9600/line speed} D->O: <75,...,48000>{9600/line speed}

Fast select: {N}|Y|RespRestr Des. max. transit delay (msec): <1 to (2**16 - 1)>

CUD field: Initial hex: <even # of hex. digits> Last ASCII: <char. string>



Help & Screen Action Requests: /0-9,C,E,F,P,N,R, or T/

<Entry Error Diagnostics & Assistance Messages Displayed Here>

/Note 1: Don't locally echo second NUI field value/

/Note 2: CUD field cannot exceed 16 octets (128 if fast select)/

SDLC PAD CONTROL MODE

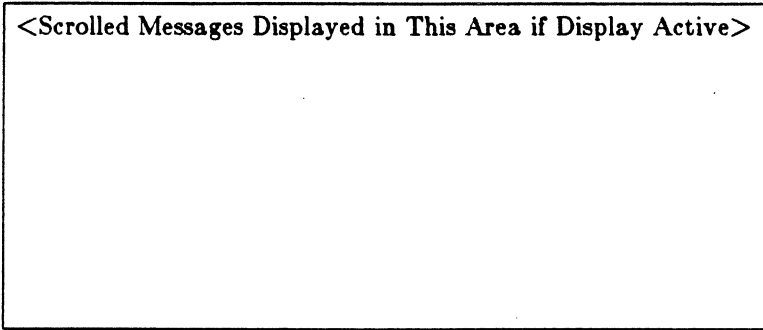
SCREEN #3: SVC CLEAR

Address of orig. party: *<PPSN addr. or abbrev.>*

Logical channel number: *<1-4095>*{*If only one LC active for addr., that LC*}

Clear user data: *<ASCII character string up to 16 char. (128 for fast select)>*

<Scrolled Messages Displayed in This Area if Display Active>



Help & Screen Action Requests: *[0-9, C, E, F, P, N, R, or T]*

<Entry Error Diagnostics & Assistance Messages Displayed Here>

BLANK PAGE

References

1. "Systems Network Architecture Concepts and Products," IBM Document GC30-3072-3, October 1986.
2. "Systems Network Architecture Technical Overview," IBM Document GC30-3073-2, September 1986.
3. "Public Packet Switched Network Generic Requirements," Bellcore Technical Reference TR-TSY-000301, Issue 1, September 1986.
4. "Proposed Revisions to the Public Packet Switched Network Generic Requirements (PPSNGR)," Bellcore Technical Advisory TA-TSY-000301, Issue 2, May 1988.
5. "Operations Technology Generic Requirements," TR-TSY-000439, Issue 2, February 1988.
6. Synchronous Data Link Control Concepts, IBM Document GA27-3093-3, June 1986.
7. "Format and Protocol Reference Manual: Architecture Logic for Type 2.1 Nodes," IBM Document SC30-3422-0, December 1986.
8. "IBM Token-Ring Network Architecture Reference, IBM Document SC30-3374.
9. "The X.25 Interface for Attaching SNA Nodes to Packet-Switched Data Networks General Information Manual," IBM Document GA27-3345-2, March 1985.
10. "Tuning Considerations for IBM SNA X.25 DTE's," IBM Document GG24-1746-00, October 1985.
11. "X.25 NCP Packet Switching Interface for the IBM 3725 and 3720, Installation and Operation," IBM Document SC30-3201-4, December 1986.

NOTE

All Bellcore documents are subject to change and their citation in this document reflects the most current information available at the time of printing. Readers are advised to check current status and availability of all documents.

Technical Advisories (TAs) are documents that describe Bellcore's preliminary view of proposed generic requirements. To obtain TAs, write to:

Bellcore
Document Registrar
445 South Street, Room 2J125
Box 1910
Morristown, NJ 07960-1910

To obtain other Bellcore documents, contact:

Bellcore
Customer Service
60 New England Avenue, Room 1B252
Piscataway, NJ 08854-4196
(201) 699-5800

Region/BOC personnel should contact their company document coordinator and Bellcore employees should call (201) 699-5802 to obtain documents.