

Bellcore

Ⓐ Bell Communications Research

Technical Reference
TR-TSY-000885
Issue 1, December 1989

PPSN Support of SNA/SDLC Interfaces

This document replaces:

PPSN Support of SNA/SDLC Interfaces
TA-TSY-000885, Issue 1, November 1988.

This document cannot be reproduced
without the express permission of Bellcore,
and any reproduction, without authorization,
is an infringement of Bellcore's copyright.

Copyright ©1989,
Bellcore
All rights reserved.

TECHNICAL REFERENCE NOTICE OF DISCLAIMER

This Technical Reference is published by Bell Communications Research, Inc. (Bellcore) to inform the industry of Bellcore's view of the proposed generic requirements for *PPSN Support of SNA/SDLC Interfaces*.

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, EXPRESS 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 or any Bellcore Client Company (BCC)* to purchase any product whether or not it provides the described characteristics.

Readers are specifically advised that each BCC may have requirements or specifications different from the generic descriptions herein. Therefore, any vendors or manufacturers of products should communicate directly with a BCC 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:

C. F. Klein, Jr.
District Manager, Packet Switching Technology
Bellcore
331 Newman Springs Road, Room 2X309
Red Bank, NJ 07701-7030

For general information, please contact:

District Manager
Information Exchange Management
Bellcore
P.O. Box 1910
445 South Street, Room 2K-122
Morristown, NJ 07960-1910

* Bellcore Client Company (BCC), as used in this document, means any divested Bell Operating Company, or its successor, or any regional affiliate thereof.

PPSN Support of SNA/SDLC Interfaces

CONTENTS

1. INTRODUCTION	1-1
1.1 Relation to TA-TSY-000885	1-1
1.2 Background on SNA and SDLC	1-1
1.3 Current IBM Support of X.25 Interworking	1-2
1.4 Support of SNA/SDLC and X.25 DTEs	1-2
1.5 Network Capabilities Under Consideration	1-3
1.6 Terminology	1-4
2. GENERAL REQUIREMENTS	2-1
2.1 Types of Connections Supported	2-1
2.2 Options Defining PAD Types: TPAD, HPAD, and SPAD	2-1
2.3 Integrating PAD Functions into the Network	2-2
2.4 Compatibility with SNA/SDLC Products	2-3
2.5 Support of Customer Network Management Tools	2-3
2.6 Billing	2-4
3. PAD SDLC INTERFACE SUPPORT	3-1
3.1 Physical Layer	3-1
3.2 SDLC Link Layer	3-2
3.2.1 SDLC Protocol	3-2
3.2.2 Logical Link Control	3-5
3.2.2.1 QLLC Support	3-5
3.2.2.2 PAD Polling Functions	3-7
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 and PUs to and from VCs	5-1
5.2 X.121 Address Assignments and Hunt Groups	5-3
5.3 Interpreting and Mapping XID Parameter Values	5-3
5.4 Virtual Circuit Setup, Clearing, and Resetting	5-5
5.4.1 Permanent Virtual Circuits	5-5
5.4.2 Switched Virtual Circuits	5-6
5.4.2.1 All Virtual Calls	5-6
5.4.2.2 PAD-Initiated Calls	5-10
5.4.2.3 DTE-Initiated Calls	5-11
5.5 Virtual Circuit Profiles	5-12
5.6 Data Transfer Phase Functions	5-14

5.7 Security Features	5-14
5.7.1 Access Control	5-14
5.7.2 Disconnection Protection	5-15
6. SDLC/VC MAPPINGS AND STATE CHANGES	6-1
6.1 Mapping SDLC Frames to/from Packets	6-1
6.2 Send Queues and Flow Control: SDLC and Virtual Circuit	6-2
6.3 Relationships Between SDLC Interface and Corresponding VCs	6-3
7. GLOSSARY	7-1
7.1 Definition of Terms	7-1
7.2 Acronyms	7-3
8. REFERENCES	8-1
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 and Display Screens	E-1

LIST OF FIGURES

Figure 1-1. PPSN Support of IBM SNA Products Offering X.25 Interfaces	1-5
Figure 1-2. PPSN Support of Generic SNA Products Offering X.25 Interfaces	1-6
Figure 1-3. PPSN Supports SNA/SDLC Products Directly	1-7
Figure 3-1. SDL Flowchart of SDLC PAD Polling of Secondary Stations (Sheet 1 of 2)	3-11
Figure 5-1. Illustrating SNA Support of Simultaneous Access to Different Hosts	5-17
Figure 5-2. Normal Chronological Relationship Among Physical, Link, and SVC Layer Events	5-18
Figure 5-3. Time Relationships Among Link Layer and SDLC Establishment Events (Sheet 1 of 4)	5-19
Figure A-1. Illustration of SNA Services Processor Concept for LU Switching	A-3
Figure E-1 Screen #0 Display Format	E-2
Figure E-2 Screen #1 Display Format	E-3
Figure E-3 Screen #2 Display Format	E-4
Figure E-4 Screen #3 Display Format	E-5

LIST OF TABLES

Table 2-1. Protocol Type and PPSN Connection Support Summary	2-5
Table 2-2. Characteristics of Access Interface Types on PPSN	2-6
Table 3-1. Summary of Physical and Link Layers of PAD SDLC Interface Support	3-10
Table 4-1. SDLC PAD Support of SVCs and PVCs (Sheet 1 of 2)	4-4
Table 5-1. Valid SDLC PAD Interface Configuration Combinations	5-23
Table 5-2. Two-Digit (Partial) VC Profiles With a Uniform Definition Across PPSNs	5-23
Table 6-1. QLLC Mapping: From Local SDLC Interface Frames	6-7
Table 6-2. QLLC Mapping: From Virtual Circuit Qualified Data Packets	6-8
Table 6-3. SDLC Frames: Summary of PAD Send Conditions & Actions on Receipt (Sheet 1 of 3)	6-9
Table 6-4. QLLC Packets: Summary of PAD Send Conditions & Actions on Receipt (Sheet 1 of 2)	6-12

1. INTRODUCTION

This Technical Reference (TR) presents Bellcore's view of proposed generic requirements for a Synchronous Data Link Control (SDLC) Packet Assembler/Disassembler (PAD) function* within the Public Packet Switched Networks (PPSNs) operated by Bellcore Client Companies (BCCs). Additionally, this TR discusses PPSN support of X.25 products in a Systems Network Architecture (SNA) environment.

An SDLC PAD permits end user equipment with an SNA/SDLC interface to communicate with remote equipment via Virtual Circuits (VCs) supported by the PPSN. The SDLC PAD converts the native SNA/SDLC protocol to one that is capable of being carried over a packet network VC. 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 VC connections between two SNA-compatible devices.

The requirements proposed herein are ultimately intended to be incorporated into a future version of the PPSN Generic Requirements (PPSNGR), TR-TSY-000301^[1]; all PPSNGR references in this document refer to the PPSNGR, Issue 2, December 1988.

1.1 Relation to TA-TSY-000885

This TR represents a major revision of TA-TSY-000885, Issue 1, which it replaces. The revisions were made in response to industry comments received on the TA. Among the changes made were the following:

- Specifications for support of VC connections that do not carry the Qualified Logical Link Control (QLLC) protocol (i.e., "type B" connections) were deleted.
- The distinction between "fixed Physical Unit/Logical Unit (PU/LU) characteristics" and "grouped PU/LU characteristics" SDLC ports was removed.
- Modifications were made to the format and use of the Call User Data (CUD) field for Call Request packets.

1.2 Background on SNA and SDLC

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*^[2]

* This function would most often be implemented within the Access Concentrator (AC) network element, but could also be implemented within a Packet Switch (PS).

- *Systems Network Architecture Technical Overview.*^[3]

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.

1.3 Current IBM Support of X.25 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 International Telegraph and Telephone Consultative Committee [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 [FEPs]). 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 VC 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 IBM's 5973-L02 Network Interface Adapters (Network Interface Adapters [NIAs] are no longer actively supported).

1.4 Support of SNA/SDLC and X.25 DTEs

There are two basic levels of support that a public packet network such as a BCC PPSN could provide for SNA/X.25 interworking. In both cases, the interworking consists of using VC connections in place of physical SDLC links between SNA-compatible equipment. The two levels are:

1. Offer VC 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 ensure 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 Customer Premises Equipment (CPE) provided by any vendor.

2. A more ambitious level of support is to also provide protocol conversion support as a network service, so 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 VC 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), TR-TSY-000301^[1] 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.

1.5 Network Capabilities 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 VC 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 (X.25 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:

- System Services Control Point(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 VCs)
- 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 Binary Synchronous Communication, or 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 BCC 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 require modification over time to reflect SNA changes and service improvements.

Appendix A presents a high-level description of how such processors might be used. It suggests a possible approach to illustrate the concept, and contrasts it with the alternative of implementing such functions within each SDLC PAD.

1.6 Terminology

The use and interpretation of terms in this document are consistent with the PPSNGR. Since this TR constitutes an extension to the PPSNGR, it is assumed that the reader is already familiar with that document.

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, Section 7 provides a list of acronyms used in this document and a brief description of key SNA terms.

Within this TR, 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" is used to describe items that are subject to further study by CCITT.

Although specific details are not yet available for "under consideration" items, the label is a useful alerting mechanism that 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).

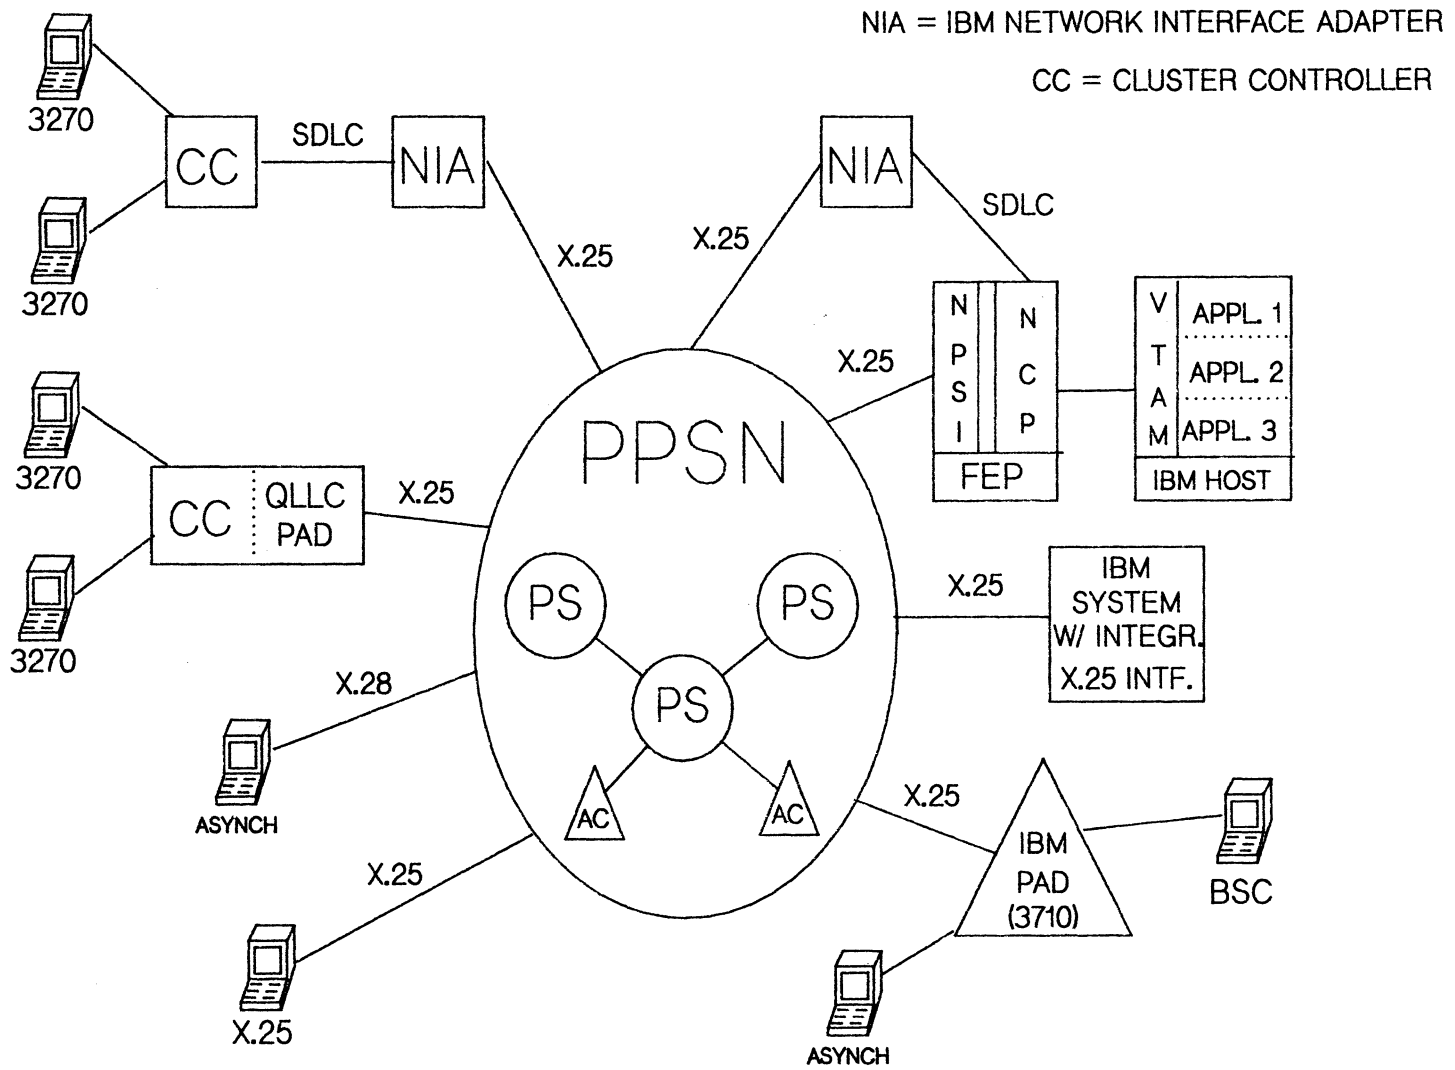


Figure 1-1. PPSN Support of IBM SNA Products Offering X.25 Interfaces

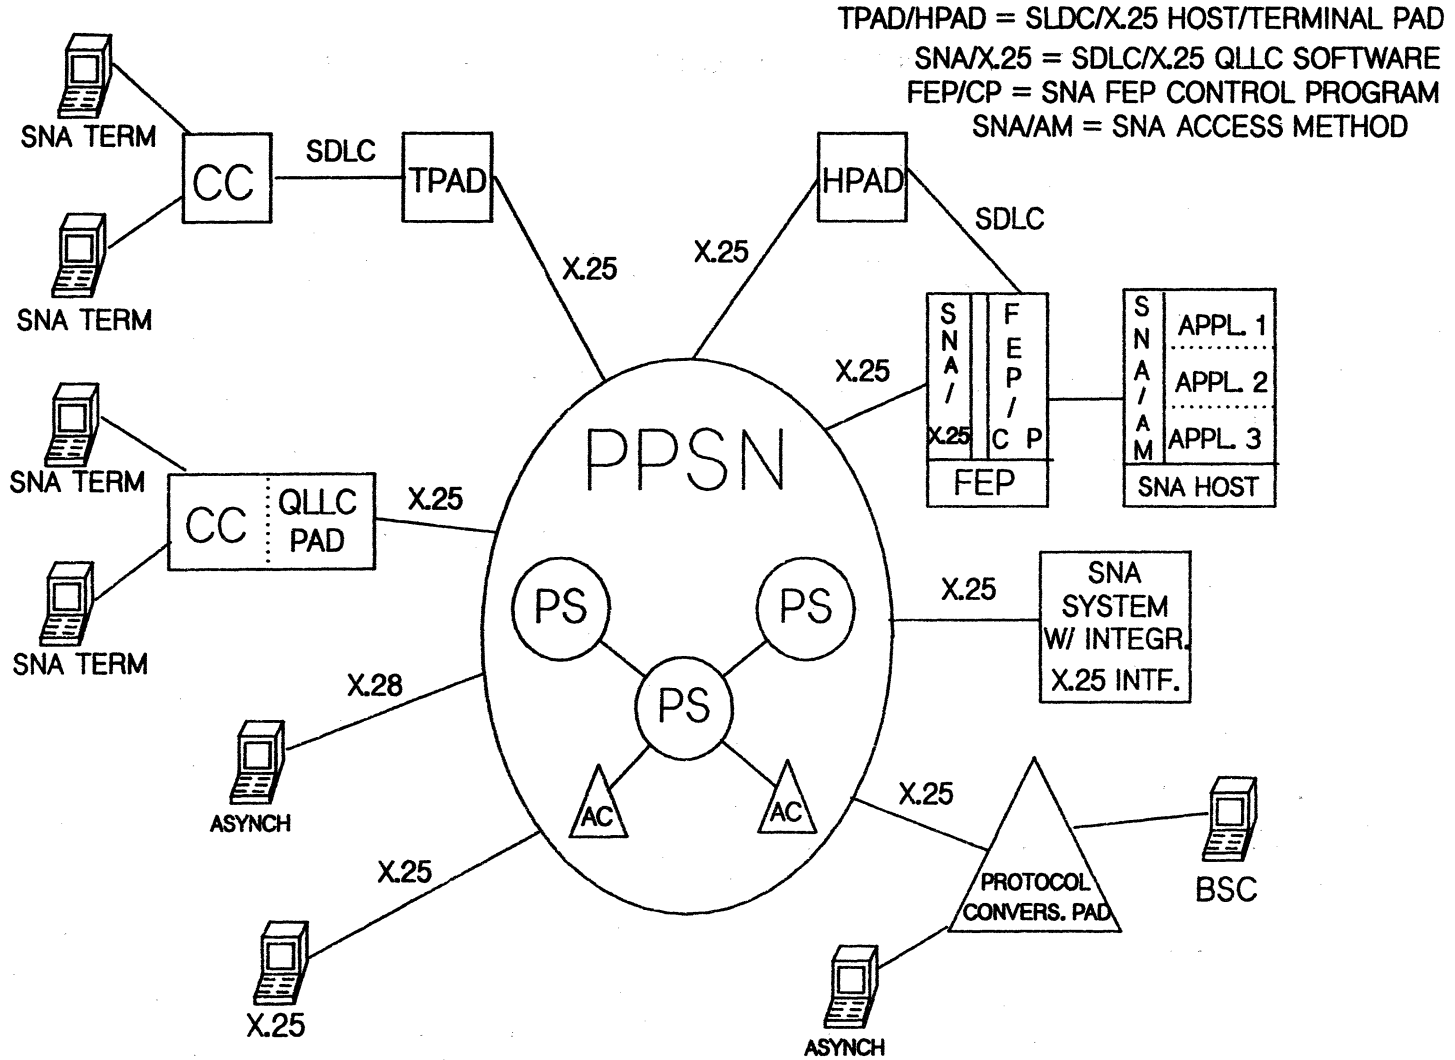
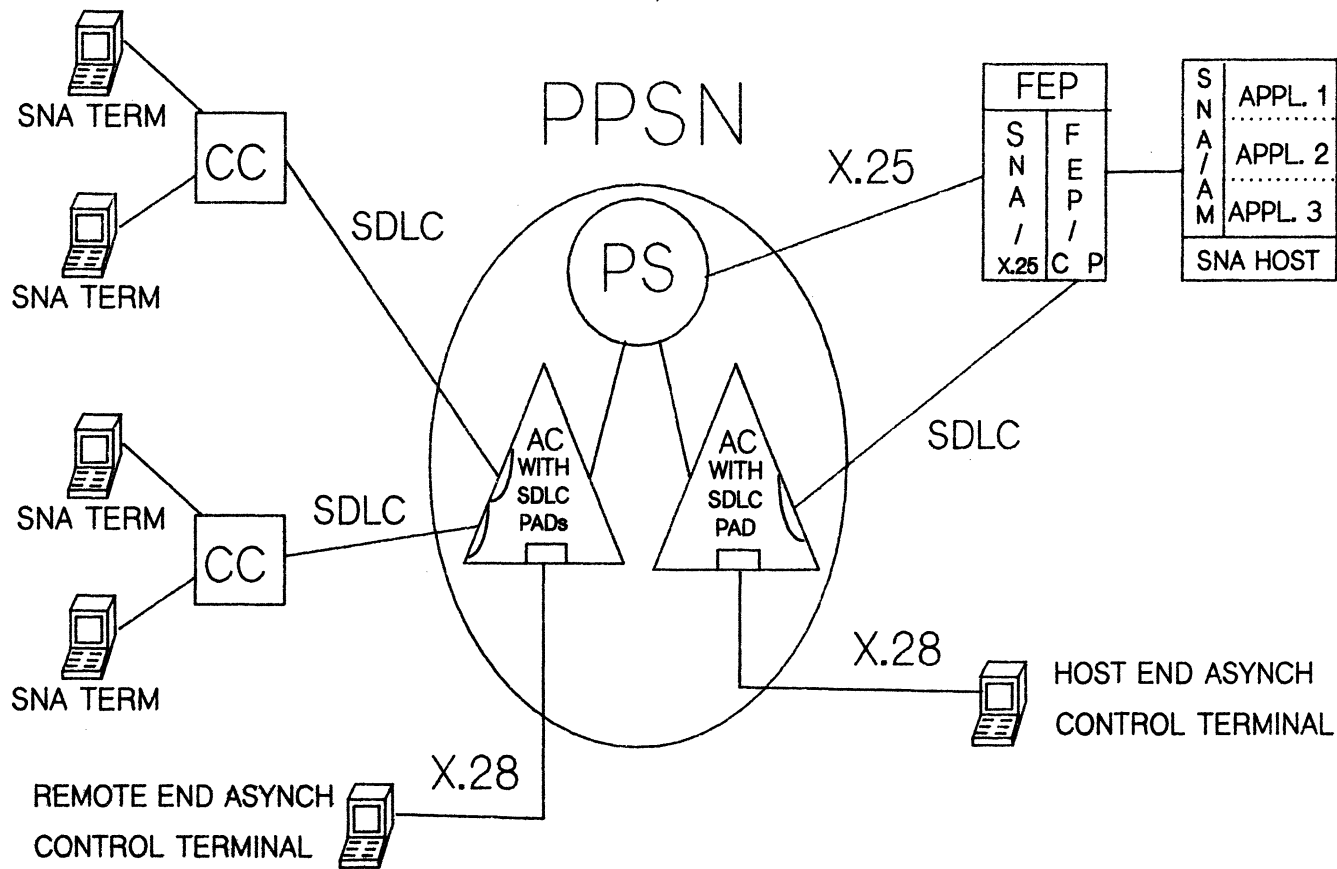


Figure 1-2. PPSN Support of Generic SNA Products Offering X.25 Interfaces

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



(SUPPORT OF DIAL-UP THROUGH PSTN NOT EXPLICITLY SHOWN)

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

2. GENERAL REQUIREMENTS

This section presents general requirements for the SDLC PAD.

2.1 Types of Connections Supported

The types of connection that are addressed by this document are characterized as follows: the DTE on each end of the VC connection is either an SNA SDLC DTE or an X.25 DTE using the IBM QLLC logical link control protocol. Both ends may be SNA SDLC DTEs, both may be X.25 DTEs using QLLC, or one of each type may be present for the connection. If the DTE is an SNA SDLC device, it is connected to the virtual circuit connection via an SDLC PAD (the requirements for which is the focus of this document). The SDLC PAD converts between SDLC and the native packet protocol of the network for VC 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 layer connection between the two DTEs, in spite of the intervening packet network(s). If both ends are X.25 DTEs using QLLC, no SDLC PADs are necessary, the QLLC passes transparently over the VC connection, and from the perspective of the PPSN, this case degenerates into a simple X.25 connection.

The mode of operation of an SDLC PAD is determined by the SDLC station type (primary, secondary, or peer) of the DTE to which the PAD is attached. The difference in PAD function associated with the type of SDLC device to which it is attached is 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. The protocol(s) that are passed within the packet network(s) are indicated for each connection supported by PPSN. The first interface column and first interface row of the table cover the 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 column or row, other types of connections involving SNA or SNA-compatible DTEs are possible across a packet network. (The fourth row and column correspond to direct X.25 access from a DTE supporting the QLLC protocol.) However, such connections are either handled by the network in accordance with the X.25, asynchronous, or 3270 BSC Display System Protocol (DSP) interface specifications of PPSNGR Sections 3.2, 3.3, or 3.8, respectively, or are not directly supported on PPSN. When an IBM protocol is involved (e.g., Enhanced LLC [ELLC] or Physical Service Header [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 CUD of the Call Request packet (see Sections 3.2.2 and 5.4.2.1). In some cases, an SDLC PAD may append details on local PU type and other information 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 that act as secondary (slave) stations, such as cluster controllers and terminal devices. The Host PAD (HPAD) connects to equipment that acts as primary (master) stations, such as mainframe computer FEPs. This distinction has been useful because of the differences in VC establishment, polling, flow control, and frame handling procedures that apply to PADs servicing the two link layer station roles, respectively.

However, recent extensions to IBM SNA allow for peer-to-peer communications in addition to that for subarea links such as those between node type 4 communication controllers. Specifically, for communication between nodes of type 2.1, the primary and secondary roles can be negotiated. To reflect the need to support negotiable role link stations 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 TR and the term Symmetric PAD (SPAD) is introduced. 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, some PAD parameters may be dynamically altered, based on the type of device on the remote end of the VC 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, secondary, or negotiable station (for which the primary/secondary role may be negotiated or determined dynamically, based on the other station with which it is communicating). The PAD would have configuration options controlling the means available for establishing VCs 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 inband SDLC terminal messages (under consideration).

In terms of these PAD options, the traditional SDLC TPAD would correspond to the secondary station DTE end. A traditional HPAD would correspond to the primary station DTE end. A PAD configuration capable of supporting IBM SNA node type 2.1, 4, or 5 devices for peer-to-peer connections would correspond to a combined station. Such a PAD is labeled a SPAD for convenience.

2.3 Integrating PAD Functions into the Network

The following requirements apply to the integration of PAD functions into the PPSN:

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 BCC customer demand at the location served by the AC.
2. Additionally, a supplier may optionally support SDLC PAD ports within PSs.
3. Common AC and PS functions such as routing and Operations, Administration, and Maintenance (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*^[4] (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 VC 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 is carried over this VC connection. The actual internodal and internetwork protocol(s) supporting the X.25 VC 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 VC service on an end-to-end basis. Interfaces between nodes of different manufactures should be X.75' within the BCC network and X.75 between networks.

2.4 Compatibility with SNA/SDLC Products

The following requirements pertain to SDLC PAD compatibility with SNA/SDLC products:

1. The SDLC PAD should be compatible with SDLC interfaces on all type 1 and 2 SNA nodes at the secondary station end, and on 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 and 4 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	

2. 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

Various 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 not interfere with the transparent operation of such network management tools.

NetView is a trademark of International Business Machines Corporation.

When an SDLC interface is replaced by a 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 (AFMs) 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 passively monitor the SDLC access interface and/or transmit 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 and 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

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

Legend

NS Not supported on PPSN
X.25+ X.25 with SNA protocol ID in CUD; see Section 5.4.2.1
QLLC+ QLLC, possibly supplemented with SDLC PAD information in CUD (see Section 5.4.2.1)
Note 1 NPSI connection to a DSP PAD would require special user software

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

Access Type	PPSNGR Section	PPSN Support	SNA Category	QLLC LLC Type
SDLC PAD	3.9*	Req.	SNA	LLC 3
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
- NA Not Applicable
- * Assumes this document may be incorporated into a future issue of the PPSNGR as Section 3.9.
- ** 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.

3. PAD SDLC INTERFACE SUPPORT

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

3.1 Physical Layer

In this document, the term “dial-up” refers to either dial-in, dial-out, or both. The term “dedicated” refers to direct connections between the DTE and the PAD that are not dial-up through the Public Switched Telephone Network (PSTN).

The following requirements apply:

1. All SDLC PADs should support dedicated interfaces.
2. Support of dial-out (i.e., from the PPSN to the DTE via switched connections through the PSTN) by SDLC interfaces configured for support of SVCs is desirable. Configuration options for dial-out interfaces are set by service order. The dial-out interface used to complete a dial-out is determined via the registered X.121 address designating the called party (see Section 5.2).
3. The SDLC PAD should support private dial-in interfaces configured with auto call specifications for completing a SVC. Configuration options for private dial-in interfaces are set by service order. See Section 5.4.2.2 for additional details on auto call service.
4. Support of public dial-in interfaces may be needed for ports optioned for per-call establishment of SVCs inband over the SDLC interface. Requirements for SDLC public dial-in and support of inband SVC establishment over an SDLC interface are under consideration.
5. Line speeds of 1200, 2400, 4800, 9600, 19200, and 56000 bps should be supported on dedicated interfaces. Line speeds of 1200, 2400, 4800, and 9600 bps should be supported on dial-up interfaces. Additional support of other line speeds is desirable.
6. The interface should conform to EIA 232D 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.
7. For 1200-bps dial-up interfaces, the PAD's modem should be EIA 232D compatible and conform to the Bell 212A-type data set. For 2400-bps dial-up interfaces, the PAD's modem should be EIA 232D compatible and conform to CCITT Recommendation V.22*bis*. For dial-out, V.22*bis* should be the default modem type.
8. For 4800-bps and 9600-bps dial-up interfaces, the PAD's modem should be EIA 232D compatible and conform to CCITT Recommendation V.32.
9. The automatic answering procedures for dial-in interfaces should be compatible with CCITT Recommendation V.25. If the PAD and associated modem functions are not physically integrated, procedures for signaling the called number on dial-out interfaces should be compatible with CCITT Recommendation V.25*bis*.
10. Dial-out from an SDLC PAD should be based on X.121 registered addresses. See Section 5.2 for details.
11. For dial-up 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 VC(s), and disconnect and deactivate the link layer associated with the

dial connection, if still active. 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., Exchange Station Identification [XID] frame) normally required for new physical connections.

If the modem function is not physically integrated into the PAD, the modem should remain "off hook" with respect to the dial-up connection until the EIA 232D CD circuit is placed in the "ON" condition by the PAD. The PAD should not place the EIA 232D CD circuit in the "ON" condition until all the VCs associated with the broken connection have been cleared. This prevents the analog switch from perceiving the analog line as being available to accept an incoming PSTN call when the PAD is not yet ready for a new incoming call (e.g., if the line is part of a hunt group, it is more efficient for the line to be passed over for a truly available port rather than permitting a "ring, no answer" situation).

12. For dial-up interfaces, 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.
13. Both duplex and half-duplex operation should be supported as interface configuration options.
14. Refer to Section 6.3 for requirements on PAD responsibilities in maintaining the physical layer and relationships between the local SDLC interface and associated VC connections.

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 (nonbasic) 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 VCs per PU are under consideration.

3.2.1 SDLC Protocol

The SDLC DTE/DCE interfaces supported by the SDLC PAD should conform to the IBM document, *Synchronous Data Link Control Concepts*.^[5]

In support of communication between SNA nodes capable of peer-to-peer communication, the SDLC interface should accommodate the ability to negotiate or dynamically set the role of primary or secondary station. Discussion of configurable (i.e., node type 4 or 5) link stations and XID exchanges used to set primary and secondary role for each such station are discussed in the IBM document, *Systems Network Architecture Format and Protocol Reference Manual: Architecture Logic*.^[6] For type 2.1 nodes, XID/XID3 formats, and associated role negotiation procedures, additional details are available in a separate IBM document, *Format and Protocol Reference Manual: Architecture Logic for Type 2.1 Nodes*.^[7]

Because of the mediating influence of SDLC PADs, local characteristics of SDLC interfaces and stations at either end of a VC connection need not match to the degree necessary when the SDLC DTEs are directly connected. For example, the SDLC stations at opposite ends of a VC connection supported by two SDLC PADs may be operating with different link layer modulus and maximum frame sizes.

The following requirements apply:

1. Link layer procedures and all frame formats specified in the IBM document on SDLC^[5] are to be supported by the PAD, unless otherwise stated below.
2. The SDLC PAD should not generate a REJ frame, but should be able to correctly respond locally to a valid REJ frame received over the local SDLC interface. (The PAD may support the ability to send REJ frames if this ability is implemented as a nondefault configuration option.)
3. It should be possible to configure the SDLC interface for either point-to-point or multipoint operation. In multipoint operation, the SDLC PAD should be able to accommodate up to 2 stations for line speeds of 2400 bps or lower, up to 8 stations for line speeds of 4800 to 9600 bps, and up to 16 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. For multipoint operation, it should be possible to assign any valid station address to each configured secondary station. For point-to-point operation, it is desirable that the single secondary station be capable of being assigned any valid station address.
4. Two-way alternate (only) operation and two-way simultaneous operation should be supported as interface configuration options. However, two-way simultaneous operation is not possible over half duplex physical layer facilities. On interfaces configured as two-way simultaneous and either dial-up or negotiable role (see below), it should be possible to override the configuration and select two-way alternate operation via the XID procedures detailed later in this document. Overriding an interface configuration of two-way alternate via XID procedures should not be permitted.
5. The SDLC PAD should support connection to a primary station, secondary station(s), 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 or on the basis of which station transmits the SNRM(E) frame. (Also see item 4 of Section 5.4.2.1 for a mechanism to pass the station role of the calling party to the PAD in the CUD of a Call Request packet.) In multipoint operation, it is desirable for a PAD locally connected to the primary station to also permit one or more secondary stations to be present on the local SDLC interface (as well as at the remote end of the VC connection). Support of this desirable capability requires that the PAD be configured with the station address of each such local secondary station and that it ignore any frame it receives with any such address (such frames are handled exclusively on the local SDLC interface and are not mapped to or from a VC connection). If this capability is not supported, then only the primary station is permitted on the local SDLC interface when the PAD interface is configured for the primary end.
6. The maximum frame size supported on the interface should be settable at least to values in the range that would accommodate information format frames with information fields of 256 to 265 octets. Support of values in the range that would accommodate information format frames with information fields of 512 to 521 octets, in addition, is desirable. For dial-up interfaces and those configured as negotiable role, it should be possible to set this value to the lesser of the frame size configured for the interface; this value should be selected dynamically through the XID procedures detailed later.
7. The SDLC PAD should support modulo 8 link layer sequence numbers and the SNRM frame type. It is desirable that the SDLC PAD support modulo 128 link layer sequence numbers and the Set Normal Response Mode Extended (SNRME) frame type. If both modulus are supported, a set modulo for the local SDLC interface at the secondary station end and a default/maximum modulo at the primary/negotiable end is set as a configuration option. (Optionally, the value may be set for each configured station address on the interface.) If both modulus are supported, the default value applying over the local interface at the

primary (or negotiable role) station end for a specified station address can be overridden downward by the mode setting command received (i.e., an SNRM instead of an SNRME if maximum modulo of 128 is configured). For modulo 8 link layer sequence number operation, link layer window values in the range 1 to 7 should be supported. For modulo 128 link layer sequence number operation, link layer window values in the range 1 to 60 should be supported. Support of window values in the range 61 to 127, in addition, is desirable for modulo 128 operation.

8. It should be possible to designate separate link layer acknowledgment window values in the range supported as a configuration option for each secondary station on the interface. For dial-up interfaces and those configured as negotiable role, it should be possible to set this value to the lesser of the value configured on the interface; this value should be dynamically selected via the XID procedures detailed later.
9. See Section 5.7.1 for additional interface configuration option requirements associated with access control.
10. The SDLC PAD should retain copies of frames that may need to be retransmitted, until their receipt is acknowledged.
11. 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 Unnumbered Information (UI) 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 VC 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 that is associated with more than one link station). However, the broadcast address, to which all secondary stations on an interface may respond, should be supported (see item 3 of Section 5.1).

However, if the RIM, SIM, and/or UI frames are supported, they should conform to the frame formats and link layer procedures in the IBM document on SDLC.^[5]

12. Refer to Section 6.1 for specification of the conditions under which the PAD transmits specific frames and PAD action in response to the receipt of specific frames. Refer to Section 6.3 for requirements on PAD responsibilities in maintaining the link layer and on relationships between the local SDLC interface and associated VC connections.
13. For local interfaces serving the secondary station end, the SDLC PAD should support retransmission of locally generated command frames in accordance with the IBM SDLC procedures. If a secondary station does not respond to a command poll that was locally generated by the PAD (i.e., not mapped from a corresponding QLLC packet from the remote end), that command should be retransmitted. (Retransmission of frames mapped from QLLC packets are handled end-to-end by the primary station DTE.) When retransmitting a frame, the PAD should wait T_r seconds for a response and should attempt N_r retransmissions before assuming an inactive state for the corresponding secondary station (see Sections 3.2.2.2 and 6.3). The PAD should be capable of configuring each SDLC interface serving secondary stations with a T_r that can be set to at least the values 0.5, 1, 2, 3, 4, and 5 seconds, and an N_r that can be set to at least the values 0, 1, 2, 3, 4, and 5.

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 VC, despite the presence of one or more intervening packet networks. This protocol rides above the X.25 packet layer and permits certain SDLC link layer 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 one logical link control option: support of QLLC (with additional PAD-related information appended to the call user data of the Call Request Packet following the protocol identifier).

For SVCs, the protocol identifier present in the first octet of the CUD field of the Call Request packet for SVCs specifies the type of Logical Link Control (LLC) that is to apply. 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 the interface that is the destination for an SVC, the SDLC PAD uses the protocol identifier in the CUD of the received Call Request or Incoming Call packet. The SDLC PAD assumes QLLC and inserts a protocol identifier corresponding to QLLC (LLC 3) in any SVC it initiates. A first octet value of x'C3' or x'CB' indicates QLLC; all other values identify protocols other than QLLC (see Section 5.4.2.1 for additional details).

If QLLC is supported over the VC, 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; these packets are then mapped to corresponding frames on the local SDLC interface. Details of QLLC support are presented in Sections 3.2.2.1, 6.1, and 6.3.

Polling functions are handled locally by the PAD and are not reflected across the VC connection. This reduces the overhead of the VC and improves performance. Details of how the SDLC PAD handles polling are provided in Section 3.2.2.2.

3.2.2.1 QLLC Support

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 VC 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 VC 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 according to IBM SDLC specifications. For example, a Receiver Not Ready (RNR) frame is used to flow control the local SDLC link. It may have an indirect impact on traffic over the corresponding X.25 VC, but it is not mapped into a specific X.25/QLLC packet.

The following requirements apply:

1. The PAD interface to the DTE should conform to Section 3.2.1.
2. Handling of XID frames across the local SDLC interface should conform to Section 5.3.
3. The SDLC frames that are mapped to and from QLLC qualified Data packets are: SNRM(E), Disconnect (DISC), XID, TEST, Unnumbered Acknowledgment (UA), Request Disconnect (RD), Disconnect Mode (DM), and Frame Reject (FRMR). Under certain circumstances, as detailed later in this document (see the next item and Tables 6-1 through 6-4), 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, except for these special

circumstances, 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. Using the QLLC protocol, RR frames and QRR packets cannot be used for end-to-end polling. As described in Section 3.2.2.2, polling is performed locally between an SDLC DTE and its PAD.

The information field of an I frame received from the DTE over the local SDLC link is mapped into the user data field of an unqualified Data packet (or packet sequence) on the virtual circuit corresponding to the local SDLC station address. Similarly, the user data field of incoming unqualified data packets (or packet sequences) is 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.

Mapping for the RIM, SIM, and UI frames is not specified within QLLC, and such mapping is not required of the SDLC PAD. If the PAD does not support such mapping, it could respond to any such frame received across the SDLC interface either with an FRMR (complete nonsupport) or locally over that interface in conformance with SDLC specifications (nonsupport of mapping). If the PAD does not support such mapping, it would respond to any QLLC-like packet associated with these frame types (i.e., QSIM, QRIM, and QUI as specified in Table 6-2) with a Qualified FRMR (QFRMR) across the virtual circuit over which the packet was received. If a PAD does support QLLC-like mapping for any of these frame types, the following apply to provide for compatibility in a multivendor environment:

- a. A PAD should not map an RIM or UI response frame to a corresponding QLLC-like packet over the respective VC connection unless the command counterpart (SIM or UI, respectively) was previously received as a QLLC-like packet over that VC.
- b. A PAD should not map an SIM or UI command frame to a corresponding QLLC-like packet or otherwise transmit such a packet over a virtual circuit connection unless a PAD configuration option exists and has been explicitly set to permit QLLC-like support of the SIM/RIM and/or UI frame types, respectively.
- c. If mapping is performed, the format of the QLLC-like packets generated (QSIM, QRIM, and QUI) should conform to Tables 6-1 and 6-2.

All other SDLC frames are handled locally, in conformance with the IBM SDLC specifications. Section 4 of the IBM X.25 interface general information manual,^[8] Section 8 and Appendix O of the IBM X.25 interface architecture reference,^[9] and Sections 5.1 and 6.1 of this TR 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.

The ability of a PAD to associate more than one VC with a single PU and/or map individual LUs to separate VCs (i.e., LU switching -- see Section 5.1) is under consideration. If the SDLC PAD implementation is capable of LU switching (i.e., a vendor-specific implementation), it should also be capable of being configured so that local handling of unnumbered frames needed for such support can be disabled (via a configuration option) when this LU switching capability is not used by the network or on that interface. (This disabling capability is to provide compliance with the above frame/packet mapping specifications.)

4. When the SDLC PAD is configured for a peer-to-peer connection (e.g., node type 2.1 or 4), it is necessary to take special actions associated with entry into the normal response mode. When the SDLC PAD serving the station that has been sent an SNRM(E) over the local SDLC interface (the "secondary" end) receives a UA acknowledgment, it should respond locally with an RNR to prevent the local DTE from entering data transfer before the remote end "primary" explicitly acknowledges receipt of the UA. This avoids a possible protocol

problem that applies only to peer-to-peer connections. Specifically, if the "primary" end fails to receive the UA response from the remote end PAD (e.g., because of line noise), and the "secondary" end enters the data transfer phase and starts to send data packets (since this is not precluded at the higher SNA layers for peer-to-peer communication), there may be no legal way to deliver associated I frames at the primary end. After the QUA is sent to the remote end and the "primary" station successfully receives the associated UA from its PAD, the primary station should explicitly signal readiness for data transfer with an RR or an I frame. In the case of an RR issued under these circumstances, the RR is mapped to a QRR for transmittal to the "secondary" end. On receipt of a QRR at the secondary end in this case, the PAD clears the RNR condition it had imposed on the station by transmitting a local RR. Receipt of data packets from the primary end would also clear the RNR condition set by the PAD in this situation.

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 VC(s) should be handled in accordance with the procedures of Section 6.3.

3.2.2.2 PAD Polling Functions

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(s) on the remote end of the VC connection. The PAD polls a directly connected secondary station on behalf of the primary station on the the remote end of the VC connection. None of these local polls are propagated across the VC connection. The use of local polling avoids inefficient use of VC capacity for passing polls on an end-to-end basis.

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

The following requirements apply:

1. The SDLC PAD should provide the following basic polling functions at the secondary station end:
 - a. Ensure 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 without an attempt to initially transmit I frames to reduce unnecessary link overhead. However, if the period of modified operation defined in item 15 of Section 6.3 applies, no polling is done until this period expires.
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.

3. 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.

The remainder of this subsection describes a specific polling approach that is designed to achieve the basic polling functions listed above at the secondary station end for the case of two-way alternate operation. The detailed procedures specified below illustrate the intent of the basic polling functions and represent a suggested implementation approach; however, the detailed procedures that follow are not requirements.

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 VC. These stored frames are 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. (Note that whether the poll bit is set by the primary at the beginning or end of a sequence of frames sent to a secondary station is determined by the nature of operation: two-way alternate, two-way simultaneous -- send and receive from different secondaries only, or two-way simultaneous -- send and receive from same secondary allowed.) The 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 VC. Section 6.2 addresses procedures used by the PAD to maintain send queues.

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 on its current polling status and the type of operation that applies (i.e., two-way alternate, two-way simultaneous for same secondary, or two-way simultaneous for different secondaries).

At any time, an operational station will be considered either "active" or "inactive" by the PAD. All operational stations associated with a VC are considered inactive when the SDLC interface is initialized. The inactive station poll can be configured to be either a format 0 XID or an SNRM(E), with the XID being the default. An inactive station becomes active following a valid response to a poll from the PAD within the timeout period. 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. 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 (or cannot be placed in) the Normal Response Mode (NRM). The PAD should establish an SVC if the interface is configured for auto call service and the secondary station returns to active status.

Assuming the case of two-way alternate operation, 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). In all cases, polling is 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 an XID, SNRM, or SNRME, as configured for the interface, with the poll bit set.

After the poll, the PAD awaits a valid frame sequence response. 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 an RR frame, as appropriate. If a valid response is not received from a currently active station after the specified number of retries, the station 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 are 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 polling procedure described above is illustrated in Figure 3-1 with a System Description Language (SDL) flowchart for the case of two-way alternate operation. This procedure represents one possible approach that is intended to achieve the basic polling functions cited in item 1 above.

For two-way simultaneous operation, the primary can independently send and receive at the same time; therefore, the transmit and receive operations/queues are not strictly sequential as in the case illustrated in Figure 3-1. The primary station could transmit I frames (without setting the poll bit) to one station while receiving frames from the (same or another) secondary station that has been polled.

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

Physical Layer

Line Speeds (bps)
Dedicated: 1200, 2400, 4800, 9600, 19200, 56000
Dial-up: 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 Layer

Normal Response Mode (NRM); Normal Disconnected Mode
Normal Response Mode Extended (NRME)*
Modulo 8 Frame Sequencing
Modulo 128 Frame Sequencing*
Link Layer Window Sizes: 1-7 (mod. 8); 1-60, 61-127* (mod. 128*)
Point-to-point and multipoint operation options
Two-Way Alternate (Only) & Two-Way Simultaneous Operation Options
Minimum Number of Stations Supported Per SDLC Interface
2400 bps or less: 2
4800-9600 bps: 8
greater than 9600 bps: 16
Support of primary, secondary, or negotiable local interface station options
Max. Frame Size to Accommodate I Frame Information Field of:
256 to 265 octets
512 to 521 octets*
Dynamically Settable By XID Exchange on Dial-Up Interfaces:
2-way alternate vs. 2-way simultaneous operation
primary, secondary, or negotiable local station
maximum I frame information field length
link layer window for each secondary station
All SDLC Frame Types Except
RIM, SIM (Optionally, may be supported)
UP, CFGR, BCN
UI (Optionally, may be supported)

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

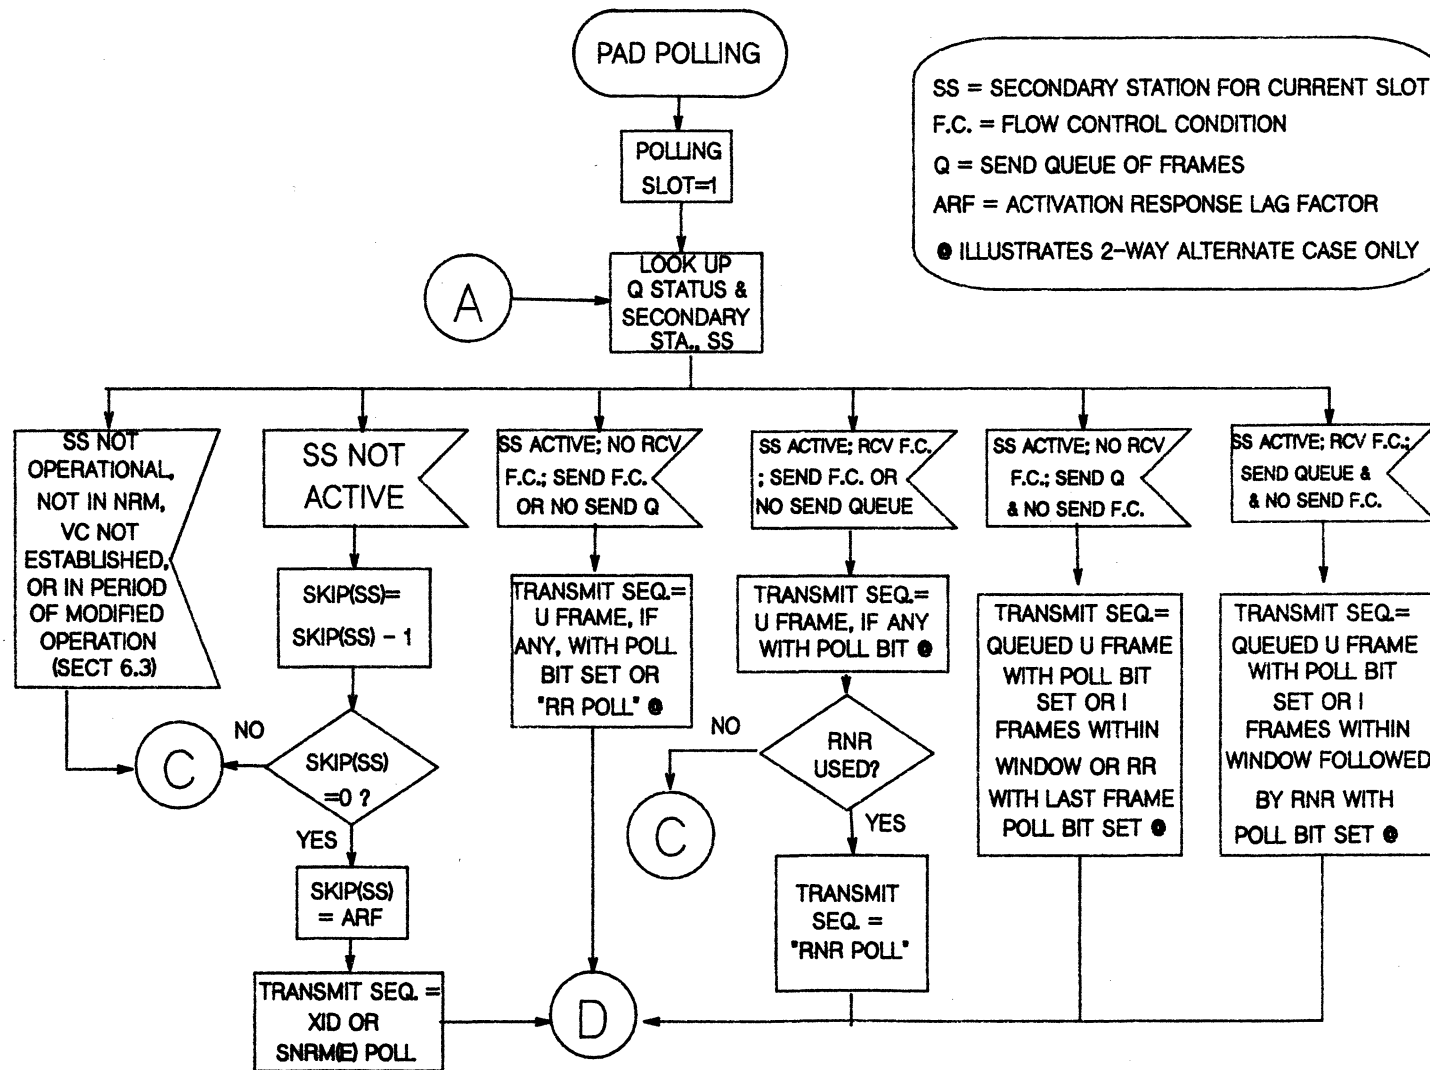


Figure 3-1. SDL Flowchart of SDLC PAD Polling of Secondary Stations (Sheet 1 of 2)

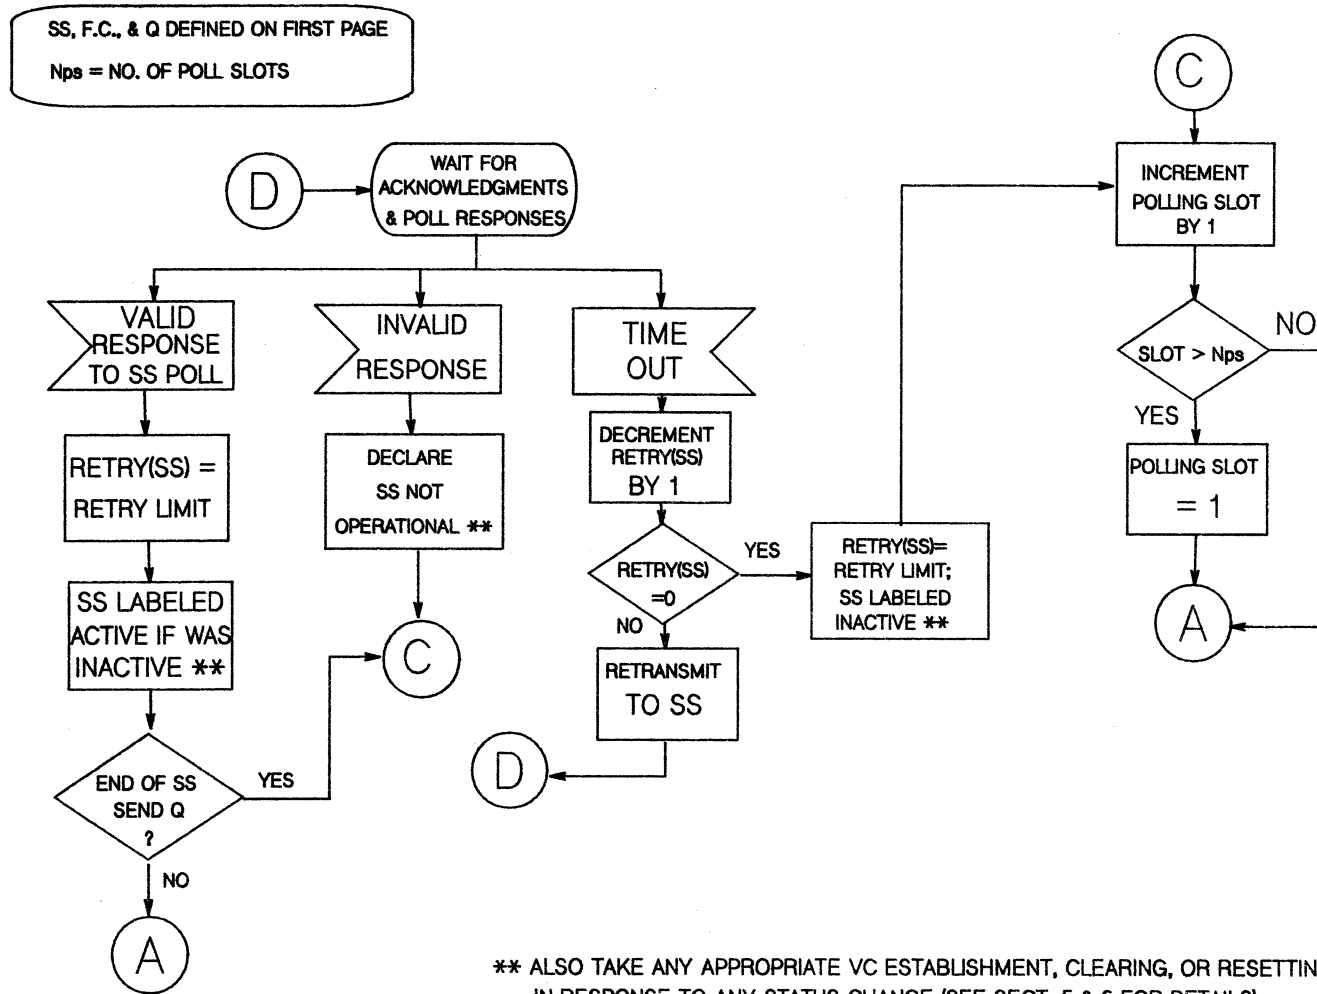


Figure 3-1. SDL Flowchart of SDLC PAD Polling of Secondary Stations (Sheet 2 of 2)

4. PAD VIRTUAL CIRCUIT SUPPORT

This section addresses SDLC PAD support of X.25 VCs.

4.1 Support of PVCs and SVCs

For each active SDLC interface that a PAD terminates locally, one or more VCs 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 BCC network of another manufacturer in establishing the VC, a logical channel on an X.75' interface between the two network elements is used. Otherwise, the SDLC PAD may maintain the VC by means of proprietary protocol interfaces with other network elements. If the remote end SDLC PAD or X.25 DTE is not on the same network, the respective VC will "leave" the network across an X.75 interface.

1. In all cases, the VC 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 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.^[1]
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	8
4800 to 9600 bps	8
over 9600 bps	16

Also see Section 5.1 for requirements on the relationship between the maximum number of secondary stations that can be configured and the number of simultaneous VCs 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 SVCs and any number of PVCs.
4. The SDLC PAD should be capable of having an SVC originate from or be destined for an SDLC interface configured at the physical layer as being either dedicated or dial-out through the PSTN. Thus, an SVC may originate from the PAD even if the associated local SDLC interface has been physically established via dial-out from the PAD (i.e., the SVC is requested by an asynchronous control interface as specified in Section 5.4.2.3).
5. 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 4 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" (x'80') and diagnostic code decimal 235.
6. Support of SVC and PVC service by the PAD should conform to the IBM QLLC specifications.^[8] ^[9] In particular, the PAD should:

- Support the Q bit mechanism for identifying QLLC commands and responses
 - Support the M bit mechanism
 - Use only the "0" value of the D bit
 - Always include the diagnostic code field in Clear Request, Reset Request, and Restart Request packets
 - Support the diagnostic codes specified in Appendix H of the above cited IBM architecture reference^[9] 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 [e.g., when an invalid protocol identifier is detected in the CUD -- cause code x'80' 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 of the IBM specifications for SNA/X.25 DTE/DCE interfaces^[9] in the generation and handling of SVC and PVC resets.
7. 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.
8. 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 option controls whether or not the PAD should clear call requests in which the local interface is an alternate destination, following call redirection. The second option 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 VC remote from the clearing PAD should receive the clearing cause "DTE originated" (x'80') and diagnostic code decimal 228 (facility not subscribed).

4.2 X.25 Product Compatibility

The PPSN supports DTEs 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. Some of the SNA products providing such interfaces are as follows:

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 Packet Switched Public Data Network (PSPDN) packet switched connection requires that careful attention be paid to several tuning considerations. These consideration 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 its 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 acknowledgments 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 IBM's *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 (Sheet 1 of 2)

Packet Layer (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: 2
 4800 to 9600 bps: 8
 over 9600 bps: 16
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
Only 0 value of the D Bit

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/IA, CUG/OA/IA
 Up to 9600 bps SDLC interfaces: 10 CUGs/interface
 Over 9600 bps SDLC interfaces: 100 CUGs/interface
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 (Sheet 2 of 2)

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

- IC Preselection
- RPOA Selection Barred
- Automatic (Auto) Call Service (see Section 5.4.2.2)
- Virtual Circuit Profiles* (see Section 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
- 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.

5. ASSOCIATING SNA/SDLC PUs WITH VIRTUAL CIRCUITS

Given the VC service provided by SDLC PADs, as described earlier, this section addresses how sessions involving SNA PUs and LUs served off a PAD-terminated SDLC line are mapped into and carried over PPSN VCs, at the service and concept level. Details of state transitions and frame and packet mapping for the SDLC and X.25 packet layer protocols involved are covered in Section 6.

5.1 Mapping LUs and PUs to and from VCs

When connecting a local SDLC interface to the remote end PAD or DTE via a VC 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 vCs for two reasons:

- The station address distinction available in the SDLC protocol is not available in the link layer service (i.e., Link Access Procedure-Balanced [LAPB]) provided over X.25 VCs. 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 VC.

Thus, each 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 VC termination at the other end. The only question is whether the SDLC PAD should be able to support more than one VC for a given local SDLC station. Specifically, should the PAD be able to associate a group of one or more LUs on a station (PU) with the group's own VCs.

This latter 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 on SNA routing capabilities associated with LU-to-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 reconfiguration without the need for Advanced Communications Function/Network Control Program (ACF/NCP) and ACF/Virtual Telecommunications Access Method (VTAM) changes. In the interface with an SNA 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 VCs 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 VC 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 Figure 5-1, respectively).

Whether or not requirements for LU mapping should be developed remains under consideration. Two factors that argue against such requirements are:

- The significant increase in protocol and buffer complexity required for mapping LUs to VCs in the PAD
- The availability of routing capabilities within SNA to permit LUs to connect to geographically separate applications.

The following requirements apply:

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 VC connection. (At a PAD servicing secondary stations across the SDLC link, the SDLC station addresses are those of the actual PUs. At a 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.) The VC (PVC or SVC) is identified at the PAD interface by its local LCN. Packets sent across the VC 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 VC 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 VCs 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} VC connections on the PPSN associated with SDLC interfaces, in addition to VCs needed to support non-SDLC access interfaces on the PAD/AC/PS. N_{vc} is the maximum number of (actual or emulated) SDLC secondary stations that can be configured for the PAD.
5. Table 5-1 identifies the combinations of SDLC PAD interface configuration options that are valid, given current requirements.
6. 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 VC) SDLC procedures concerning the order of frame responses on the local interface should not be violated by the PAD.

For SDLC interfaces configured as "switched" in an IBM FEP, NCP/VTAM permits the PU type to be specified as "either 1 or 2" (the "either" designation is one of the options). Selecting the "either" designation would be safest at the host end unless it was certain that only secondary stations of a single PU type were to access the host port. In addition, for interfaces configured as switched in NCP/VTAM, 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 designating an interface as switched is that only one secondary station can be

supported on the interface according to IBM specifications. Thus, a PAD cannot concentrate traffic from multiple secondary PUs on SDLC ports configured as switched on an FEP at the host end.

5.2 X.121 Address Assignments and Hunt Groups

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.)

Because of the one-to-one association of actual secondary stations with PU emulations of stations at the remote (primary) end of the VC connection, if hunt groups are defined, each address must correspond to a station configuration compatible with the calling station. This compatibility factor is more critical for SDLC interfaces configured as nonswitched in NCP than for those configured as switched because NCP assumes a specific PU/LU profile for nonswitched SDLC interfaces.

To simplify the selection of an appropriate dial-out port, the called address associated with dial-out to an SDLC device should always correspond to an X.121 address registered with the PPSN. Associated with such a registered address would be an E.164 address for dial out, the device type (SDLC) for PAD function selection, modem type, and (if necessary) PSTN Interexchange Carrier (IC).

The following requirements apply:

1. Each actual or emulated secondary station on an SDLC interface on an SDLC PAD should be configured with a single X.121 DTE address conforming to Section 4 of the PPSNGR. These addresses pertain to the secondary stations actually connected to the PAD (secondary end) or being emulated by the PAD (primary end). (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.)
2. It is desirable that a collection of SDLC interface (actual or emulated) secondary stations (i.e., X.121 addresses) within the same AC or PS be able to be configured as a hunt group. If the SDLC PAD supports a hunt group capability, it should conform to the specifications in Section 3.2.6.22 of the PPSNGR.
3. DTE-specified SVC setup requests that result in dial-out from an SDLC PAD are assumed to specify an X.121 address registered with the PPSN as the called address rather than an (escaped) E.164 address directly. The registered X.121 address will correspond to the appropriate SDLC PAD, which will use the appropriate modem type and E.164 address (stored in memory) to complete the dial-out. If necessary to complete dial-out, a PSTN IC designation is also associated with the X.121 address.

5.3 Interpreting and Mapping XID Parameter Values

The XID 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 dial-up, it is typically used immediately after a dial-up 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 on a newly established dial-up connection.

Over SDLC interfaces between two negotiable role 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 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 negotiable role node station. The PAD should also be able to initiate an XID command on behalf of a primary or negotiable role 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 negotiable role 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 a secondary station(s) or negotiable role node. XID procedures and formats should be in conformance with IBM SDLC specifications^[5] 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. The SDLC PAD should be able to ascertain dynamically selected SDLC interface characteristics on the basis of XID/XID3 negotiations or responses for dial-up interfaces and those configured for negotiable role stations. Dynamically selected characteristics are as follows:
 - Primary versus secondary station
 - Two-way alternate versus two-way simultaneous operation
 - Maximum frame size
 - Link layer acknowledgment windows.

The PAD should be able to determine these characteristics for any format 0 XID currently generated by (i.e., identifying) major device types listed in Section 2.4 that generate such XID values. See Section 3.2.1 for requirements on the use of XID in setting link layer configuration values.

3. The SDLC PAD should be able to send and receive Qualified XID (QXID) qualified Data packets in conformance with the IBM QLLC specifications,^[8] ^[9] the general frame/packet mapping specifications of Section 6.1, and the specific XID specifications contained in the remaining numbered items of this section.
4. The SDLC PAD should support a configuration option for each private dial-in interface. Support of this configuration option for other dial-up SDLC interfaces is under consideration. The configuration option specifies whether or not the SDLC PAD initiates an XID exchange across the local interface immediately after the dial connection and the link layer are established. If it does, the PAD is also configured with two XID values of up to 6 octets each. (The ability to support XID values of greater than 6 octets each is desirable.) The first value is the one to be transmitted by the PAD in the XID command frame information field (using a broadcast address). This value identifies the primary station being emulated by the PAD and may be null. The second value is that of the XID frame information field expected

in response to that command. If the received XID response does not match this stored response (for the initial "n" octets of the response, where n is the length of the stored response), the dial-up connection is broken by the PAD and no end-to-end VC connection is established. The second value may be null, in which case a valid XID response with any (including null) information field is acceptable. This capability should not preclude or interfere with any XID exchange that may subsequently occur on an end-to-end basis between primary and secondary stations (i.e., via QXID packets) once the VC connection is established (see the remaining numbered items of this section).

5. If the PAD receives a valid QXID packet from the remote end, it should pass an XID frame (command or response, as appropriate) across the local SDLC interface with an information field taken from the received packet. However, the PAD may modify an information field value before passing it in an XID frame if the modification applies to a negotiable parameter and the PAD does not support, or the receiving station is not configured for, the capability or option associated with the original value.
6. If the PAD receives a valid XID frame across the local SDLC interface (other than a response to a PAD-initiated XID command as detailed in item 4 above), it should pass a QXID command or response packet (as appropriate) over the corresponding VC with an information field matching that of the received frame. However, the PAD may modify an information field value before passing it in a QXID packet if the modification applies to a negotiable parameter and the PAD does not support the capability or option associated with the original value.
7. As discussed in Section 3.2.2.2, the PAD may use an XID or a SNRM(E) as an "inactive station" poll across the local SDLC interface at the secondary station end.
8. Refer to Sections 5.7.1 and 5.4.2.1 for requirements relating to the use of XID information carried in the call user data of the Call Request packet.

5.4 Virtual Circuit Setup, Clearing, and Resetting

This section addresses, in detail, VC establishment, clearing, and resetting for PVCs and SVCs.

In addition, refer to Section 6.3 for relationships among (a) the physical and link layers of the local SDLC interface, (b) the associated VC connection, and (c) PAD actions with respect to the first two relationships. Many of the specifications in Section 6.3 directly impact PAD handling of VC establishment, clearing, and resetting.

5.4.1 Permanent Virtual Circuits

When the VC connecting two SDLC PADs or an SDLC PAD and an X.25 DTE 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 VC 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 VC 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

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 an X.25 interface at the other end. SVCs are appropriate when terminal equipment needs to connect to host equipment at geographically dispersed locations at different times, or when the duration of the connections through the PPSN is 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 link layer becoming active (following establishment of a dial-up connection, if applicable) and becoming inactive (or losing the physical layer connection), 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 VC 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 VC 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: the passing of LLC type and other SNA station details to the remote end in the 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 VC 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. In addition, a subfield of information for use by SDLC PADs will be carried in the CUD. This additional information will be used for connection security and other purposes.

The following requirements apply:

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 diagramed 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-up connections, where applicable
 - Establishment of the physical layer of the PAD SDLC interface at the calling and/or called ends
 - Establishment of the link layer of the PAD SDLC interface at the calling and/or called ends
 - 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
 - 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 on the link layer of the PAD SDLC interface at the calling and/or called ends for the relevant secondary station.
4. The SDLC PAD should be capable of generating a Call Request packet that contains a CUD field and interpreting those portions of the CUD as specified below. The presence of the CUD field is mandatory. It should consist of between one and six subfields. The following table provides an overview of the six possible subfields of the CUD.

Octet #	Subfield Contents	Presence
1	IBM Protocol ID	mand.
1+p	Added IBM LLC Information	opt.
1+p+q	SDLC PAD Header	opt.
1+p+q+r	PAD Supplement	opt.
1+p+q+r+s	XID Value	opt.
1+p+q+r+s+t	Reserved for PAD Use	opt.
1+p+q+r+s+t+u		
1+p+q+r+s+t+u+v		
y		

$p, q = 0$ or $p = 1, 0 \leq q \leq 15$
 $r, s, t, v = 0$ or $r = 1; s = 0, 1; t = 0, 1; v = 0, 1$
 $t, u = 0$ or $t = 1; u = 6$
 $(1+p+q+r+s+t+u+v) \leq y \leq 16$ (no fast select)

- a. The first subfield of the CUD is mandatory and contains the protocol identifier. This subfield consists of a single octet. In general, the two most significant bits of this octet (8 and 7) distinguish the type of protocol identifier. The interpretation of bits 8 and 7 is as follows:

8	7	Meaning/Source
0	0	CCITT X.29
0	1	Use by Network Administrations
1	0	Reserved for Use by ISO
1	1	Used by Higher Layer Protocol

IBM uses "11" values to identify the LLC protocol that applies to the SVC being established. The IBM interpretation and encoding is specified in Section 8.1.1.1 and Figures 8-3 and 8-4 of IBM's X.25 interface architecture reference.^[9] For connections involving the version of the PPSN SDLC PAD specified in the current issue of this document, the protocol identifier should always identify QLLC as the LLC protocol. QLLC is encoded as x'C3' ("11000011" with bit 8 to the left) or as x'CB' ("11001011" with bit 8 to the left). The x'C3' value indicates that standard X.25/ISO 8208 DTE diagnostic codes and the DTE-generated cause code X'00' are to be used. The x'CB' value indicates that extended SNA-specific diagnostic codes and the DTE-generated cause code x'80' are to be used.

- b. The second subfield is optional and contains additional IBM LLC specifications. If present, this subfield immediately follows the protocol identifier and should conform to Section 4 of the IBM X.25 interface general information manual,^[8] Sections 5.2.1.6 and 8 of the IBM X.25 interface architecture reference,^[9] and specifications of the PATH macro instruction and USRFILD operand of the X25OUFT macro instruction used by NPSI to populate the first part of the CUD.^[11] If present, the first octet of this subfield is the Field Format Identifier, which defines the format of the following octets of this subfield. 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.)

Values of this subfield that may optionally be present following a QLLC (x'C3' or x'CB') protocol identifier are x'010000' or, in the case of an IBM 3710, x'0100008dddwww0000000', where "ddd" is the upstream station address (decimal digits) and "www" is the PU id # (decimal digits) used to populate the XID (x'00000' if a real PU type 2 is attached to the 3710).

- c. The third subfield of the CUD is optional. It is useful and should be present only when an SDLC PAD is at the remote end of the connection and not all 16 octets of the CUD have already been used for the first two subfields. This SDLC PAD header must be present if the fourth subfield (see below) 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 semioctet is set to all 1s (x'F') as a flag to facilitate recognition of this subfield. Note that the only permitted value for the first octet of the optional second subfield is x'01', which can be readily distinguished from the x'F_' value for the first octet 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 optional fourth through sixth subfields. The two bits of "Caller" indicate the SDLC PAD role of the calling end:

- 00 = primary station (HPAD)
- 01 = secondary station (TPAD)
- 10 = negotiable or dynamically settable role (SPAD)
- 11 = unspecified station role.

The caller type information can be used by the receiving SDLC PAD to determine what SDLC PAD configuration (if any) is contacting it. This information would be available in advance of any XID exchange for configuration purposes.

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

- d. The fourth subfield of the CUD is optional and contains supplementary information to be used by PPSN SDLC PADs. This subfield, when present, consists of one octet:

Bits:	8	7	6	5	4	3	2	1
Octet	PU Type	XID?	Reserved	Reserved	Reserved	Reserved	Reserved	Cont2

The PU Type identifies the node type of the DTE originating the call and is encoded as follows: 001 = PU type 1, 010 = PU type 2, 011 = PU type 2.1, 100 = PU type 4, and 101 = PU type 5. All other values are reserved. The receiving PAD can use this information to verify that the calling PU type is compatible with the called end node type. The XID? bit indicates whether an XID value follows as the fifth subfield. If XID? = 1, the optional fifth subfield is present; otherwise, it is not. If the Cont2 bit = 1, then the sixth subfield is present; otherwise, it is not. At this time the sixth subfield is reserved for possible future use and is not specified. Thus, the Cont2 bit should be set to 0. Bits 4 through 2 of the fourth subfield are reserved and should be set to zero if this subfield is present.

- e. The fifth subfield is optional. If present, it contains the first six octets of the XID information field that identify the SDLC station originating the call. If this subfield is populated, the XID value placed in it is obtained from a value configured on the originating SDLC PAD interface for the applicable secondary station. This information can be used by the recipient SDLC PAD to validate the calling party identity prior to completion of VC establishment (see Section 5.7.1).
- f. The sixth and last subfield is reserved for possible future use and is not specified in the current version of this document.

The total of all 6 subfields cannot exceed 16 octets. No provision is made for passing any discretionary user data (i.e., to be passed transparently from calling DTE to called DTE) because of sequencing problems associated with mapping such data to the SDLC interface.

- 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 VC 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 control terminal acting on behalf of the interface from which the Call Request will actually be 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. If the SDLC PAD is implemented with a capability to automatically clear an established SVC based on a time and/or inactivity criterion, as a configuration option, the clearing cause should be "DTE originated."
- 8. If an SVC generated by the auto call capability is cleared, the SDLC PAD that placed the Call Request should increment a retry counter and reestablish the call (following the applicable T_d period specified in Section 6.3) if the retry counter has not reached the retry limit configured. Once the retry limit is exceeded, the counter is reset and no further auto call establishment attempts are made by the PAD.

However, the retry mechanism is reset and activated if one of the following events occurs:

- There is operator intervention
- A transition to an operational link layer is detected on the SDLC interface
- The affected secondary station makes a valid response to an "inactive station" poll.

The reattempt limit should be settable at least to values in the range 0 to 5. After a normal "DTE originated" clearing (cause and diagnostic codes of decimal value 0), the retry counter is reset.

9. If an SDLC PAD receives a valid incoming call request after having 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 station 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 to the operational status of the physical interface or the link layer, or during certain types of call clearing. The SVC is cleared by a transition of the physical or link layer to nonoperational status (e.g., DTE power off or dial connection termination). If the called party is an X.25 DTE, it may also explicitly clear the call with a Clear Request.

For this SVC alternative, the auto call service provided by the PAD (not to be confused with modem automatic calling procedures that apply at the physical layer) reduces the customer's action necessary to set up the VC connection to a minimum, 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 the duration of the SVC.

The following requirements apply:

1. Support of auto call service by the SDLC PAD should be provided on dedicated interfaces and for dial-in ports designated as private. Auto call 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 for all stations so configured in response to transitions in the status of the physical and/or link layer(s) of the SDLC interface. 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.
2. For each (actual or emulated) secondary station on the interface configured for auto call service, the PAD will initiate a VC if there is no established VC and the link layer has just become operational. In the case of dial-up interfaces, the PAD should attempt to achieve an

operational link layer immediately after establishment of the dial-up connection and physical layer operation.

3. The PAD will also initiate auto call service for a VC if that circuit was cleared, as specified in item 8 of Section 5.4.2.1.
4. The PAD should clear an established VC if the link layer becomes nonoperational or the dial-up connection is lost. While there is no operational link layer on the local SDLC interface, the PAD should not attempt to initiate a new virtual call via the auto call mechanism (see also Section 6.3).

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. In addition, the customer may subsequently request clearing of the SVC.

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 to originate 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. Dedicated 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.

It is under consideration whether the concept of an asynchronous control interface should be extended to the establishment of SVCs from non-SDLC devices, such as BSC DTEs (see PPSNGR, Section 3.8). Also under consideration is whether the ability to establish an SVC from an interface on an AC or PS other than that connected to the control interface should be provided. (For a multivendor environment, such a capability would require specifying procedures and means [e.g., a new X.75' utility] for signaling the request between the control interface and SVC origination network nodes in a standardized format.)

The following requirements apply:

1. Support of asynchronous interfaces that can be configured for the SDLC interface control function is desirable. If supported, the following numbered items (in this section) apply to ACs and/or PSs providing such support and asynchronous interfaces configured for the SDLC interface control function.
2. Asynchronous interfaces configured for the SDLC interface control function 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 VC connections that originate from SDLC interfaces on the same AC or PS. Basic support of the SDLC control interface capability applies to dedicated asynchronous interfaces. It is desirable for the AC and PS to be able 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 dedicated control interface. See Section 5.7.1 for relevant security procedures used for dial access.

3. Each AC or PS supporting the SDLC PAD function and asynchronous control interfaces should be capable of configuring SDLC interfaces for the verification of control interface authority. Specifically, each dedicated SDLC interface/SDLC station combination (i.e., corresponding to an X.121 address) not configured for auto-call, and each X.121 registered address associated with a dial out SDLC interface, should be capable of being associated with at least one asynchronous control entity on the same AC or PS. If such an association is configured, the designated asynchronous control entity is authorized to initiate an SVC on behalf of the associated SDLC station on that SDLC interface. A control entity is identified by, and is recognized as being authorized to initiate SVCs for an SDLC station by means of the authorization code. Authorization codes must be submitted by a user of a control interface and must match the code (one of the codes) configured for an SDLC station or registered X.121 address before that user can initiate SVC activity on behalf of that station (identified as the calling DTE address in the SVC setup request). Although it is desirable to be able to independently associate authorization codes with each station on an SDLC interface, it is acceptable if the association applies to the entire SDLC interface (and all stations/X.121 addresses on it) uniformly.

In the case of dial-up asynchronous control interfaces, the authorization code is also associated with a PSTN address that the PAD dials back to establish a secure dial-up control interface connection. Each dial-up control interface should be capable of being configured with up to 100 authorization codes and associated PSTN addresses to be used for dial-back. 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 SVCs it has initiated. 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 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 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 1-digit VC profile number (i.e., values 0 to 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 to 29 are reserved to be uniform across all PPSNs, as specified in this document and future issues of the PPSNGR (see Table 5-2). Profile numbers in the range 30 to 49 are reserved to be uniform across all ACs and PSs within a PPSN, as set by the responsible BCC. The remaining 2-digit profile numbers (50 to 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.

The following requirements apply:

1. Support of 1-digit and 2-digit VC profiles is desirable. If supported by an AC or PS, the following numbered items also apply.
2. An AC and PS should be able to store at least 5 profiles in the range 0 to 9 for each SDLC PAD control interface configured on the AC or PS. The AC and PS should be able to store at least 20 profiles that are maintained on a uniform basis throughout the network for each value in the range 10 to 49. The AC and PS should be able to store at least 10 profiles (which need not be related to profiles with the same number in other ACs and PSs) in the range 50 to 99.
3. 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 2-digit profiles) be referenced by one or more auto call service specifications (to reduce storage requirements and facilitate network database 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.
4. A partial VC profile does not contain elements of a full profile that 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 of the call user data field in hexadecimal notation, up to 16 octets, as specified in item 4 of Section 5.4.2.1.

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 ensure adequate network performance (aggregate and per-VC), mediate between incompatible frame and packet sizes, and process any interrupts or resets received. Local polling (Section 3.2.2.2), processing speed, and the buffers provided should permit the performance objectives specified in Section 9 of the PPSNGR to be achieved. When the sizes of the information fields of I frames on the local SDLC interface and the user data field of VC Data packets are 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 VC 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 VC procedures. The need to carefully establish authorization when a control interface is used to set up VCs 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 VC connection results because the presence of a VC may distort the host's perception of the end user. Specifically, if the host end SDLC interface to the SDLC PAD is configured as nonswitched in the FEP, 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-up 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) VC 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 VCs are established in support of SNA applications.

5.7.1 Access Control

Access control consists of several mechanisms to ensure that:

- A user establishing a VC connection to an SNA host is authorized to access the host port
- A user at an asynchronous control interface is authorized to establish an SVC on behalf of the specified originating party.

In the former case, these mechanisms are needed because access may appear to be dedicated to the host even if a switched connection is actually involved. The PPSN makes use of NUI, XID, CUG, authorization code, and secure dial-back mechanisms to provide adequate access control protection in these cases.

The following requirements apply:

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 layer station identification (via XID procedures) over the local SDLC interface and remotely over virtual circuit connections supporting QLLC, as specified in Section 5.3.
3. It is desirable for each SDLC PAD interface serving a primary or negotiable role station to be able to support XID value checking as a configuration option. If supported, XID value checking consists of checking that the 6-octet XID value carried in the CUD field of the incoming Call Request packet matches one of the XID values configured as being acceptable as remote end XID values for the interface or station. If the incoming XID value does not match one of the configured values on the destination SDLC interface, the PAD clears the call with cause code x'80' and diagnostic decimal 236 (connection identifier mismatch). A configured XID value of "null" would permit an incoming call with no XID value in the CUD to be accepted, even if checking were active on the interface. XID checking by the PAD would typically be useful when an XID exchange is not performed end-to-end by the SDLC stations because of a presumption of a dedicated connection, for example. If this capability is supported, it is desirable for the the PAD to be able to be configured with more than one acceptable XID value per secondary station (PU).
4. All SDLC interfaces for which the SDLC PAD is serving a primary or negotiable role 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.
5. 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 breaks the dial-in connection and initiates a dial-out to the DTE address. In making the dial-out, the PAD selects the modem type that is 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 is 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).
6. All public dial-up SDLC interfaces supported by SDLC PADs should be constrained to be dial-out to SDLC stations across the local interface. Support of private dial-in SDLC interfaces, configured for auto call service, is required. Support of other dial-up 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 connection interruption 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 layer above the link layer (i.e., SNA messages to inform the application of the interruption at the session layer).

1. Section 6.3 addresses actions taken by the SDLC PAD at the SDLC link layer when an SVC is cleared. What action, if any, an SDLC PAD locally connected to a primary station via an SDLC interface should take at higher layers of the SNA protocol stack to notify the host if the SVC is cleared (in addition to, or instead of, current link layer actions) 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.




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

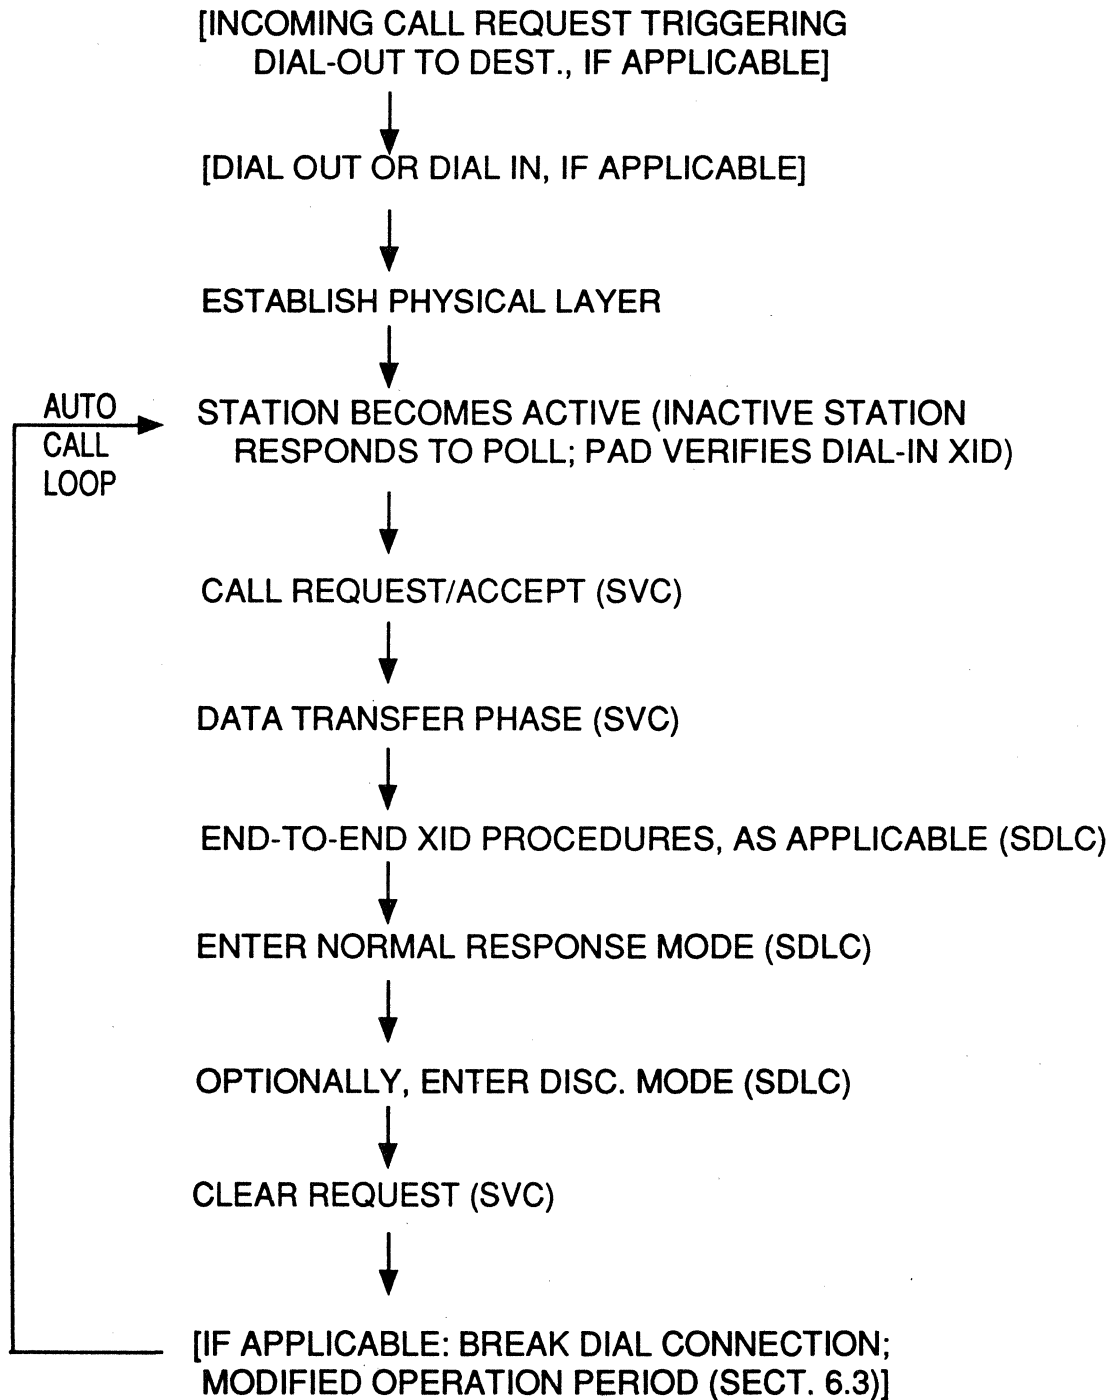


Figure 5-2. Normal Chronological Relationship Among Physical, Link, and SVC Layer Events

DTE-INITIATED CALLS: AT CALLING DTE END

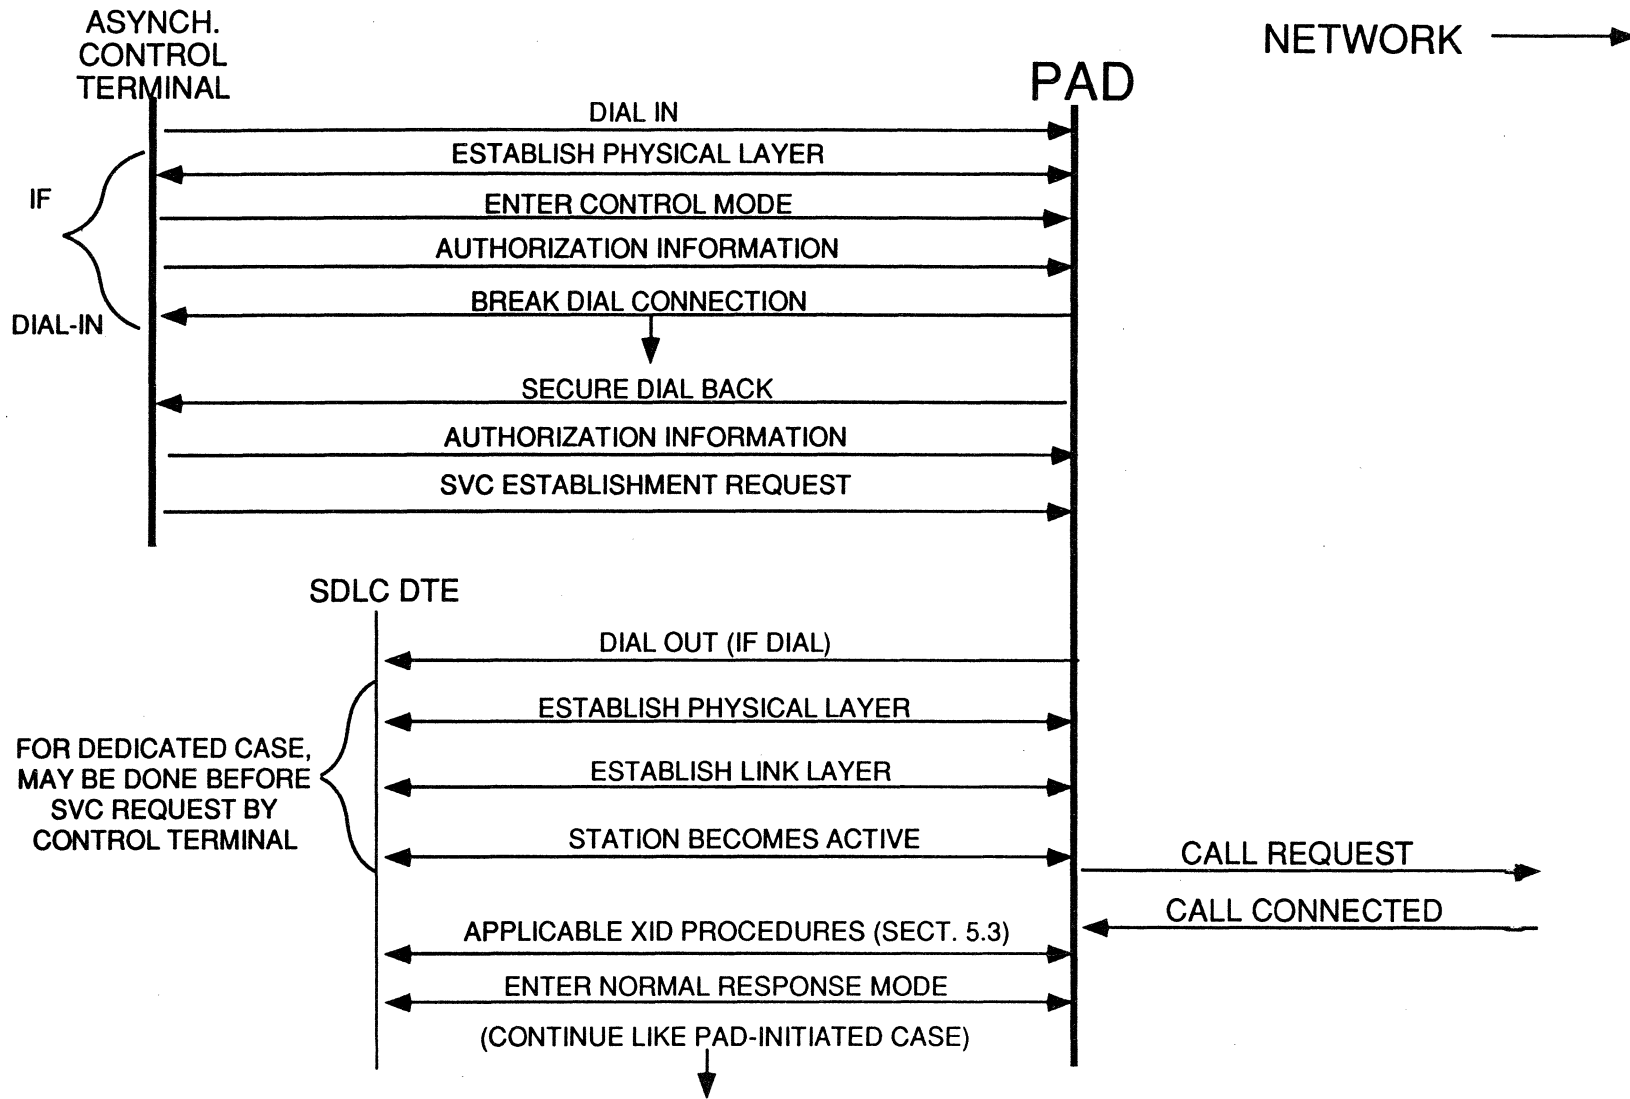


Figure 5-3. Time Relationships Among Link Layer and SDLC Establishment Events (Sheet 2 of 4)

CALLED DTE END: DEDICATED SDLC INTF.

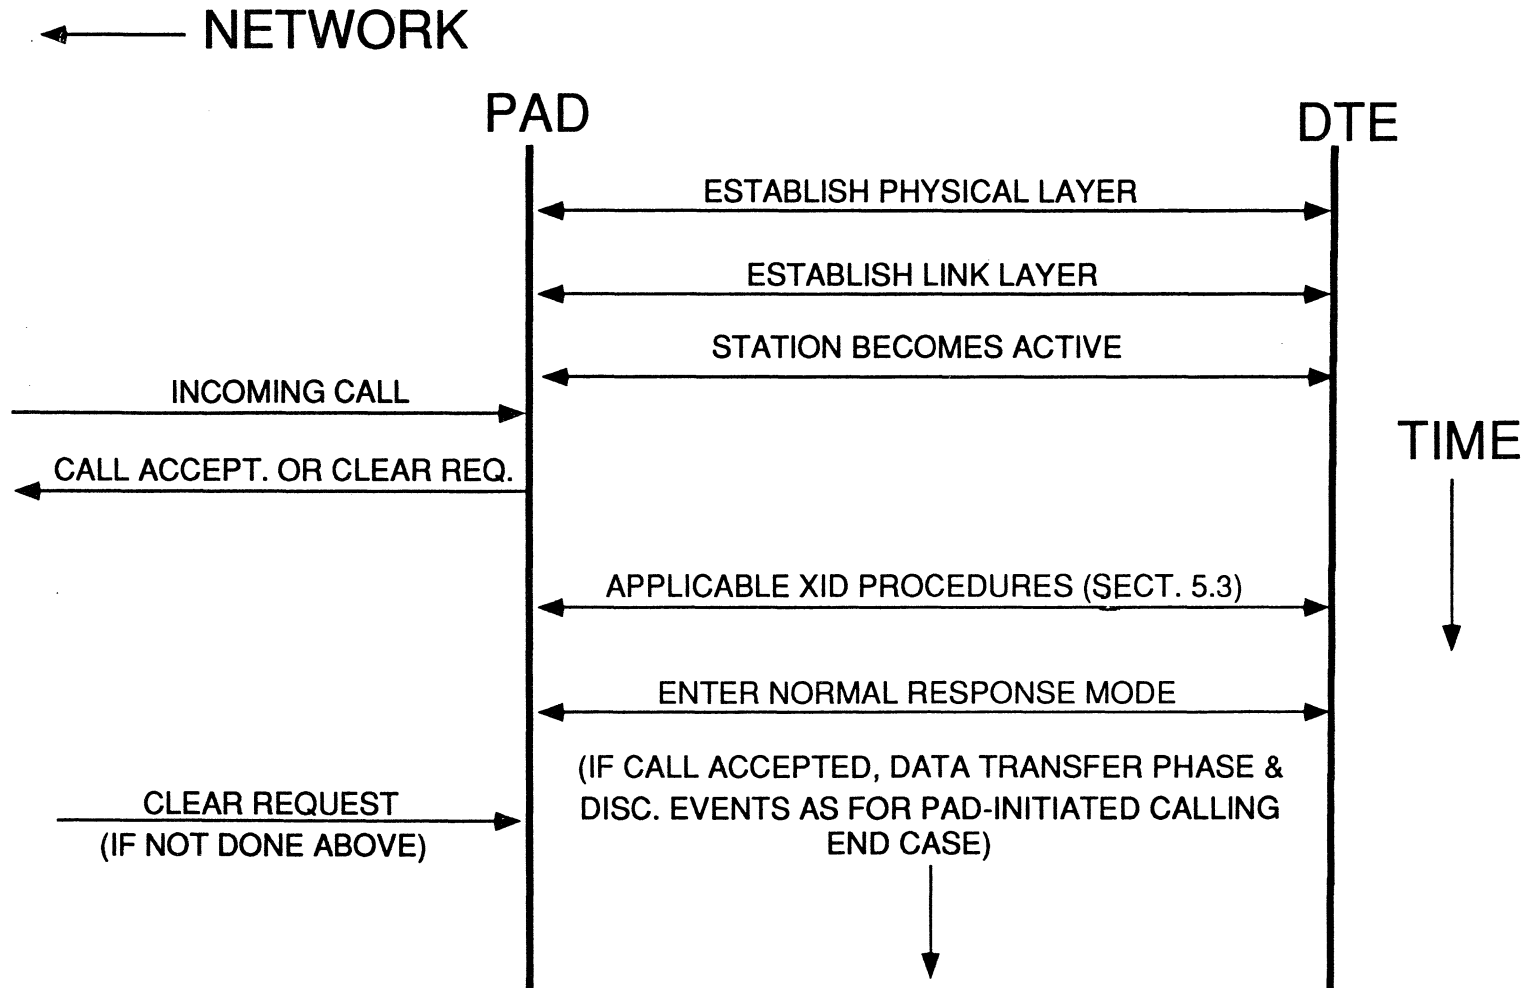


Figure 5-3. Time Relationships Among Link Layer and SDLC Establishment Events (Sheet 3 of 4)

CALLED DTE END: DIAL-OUT SDLC INTF

5-22

TR-TSY-000885
Issue 1, December 1989

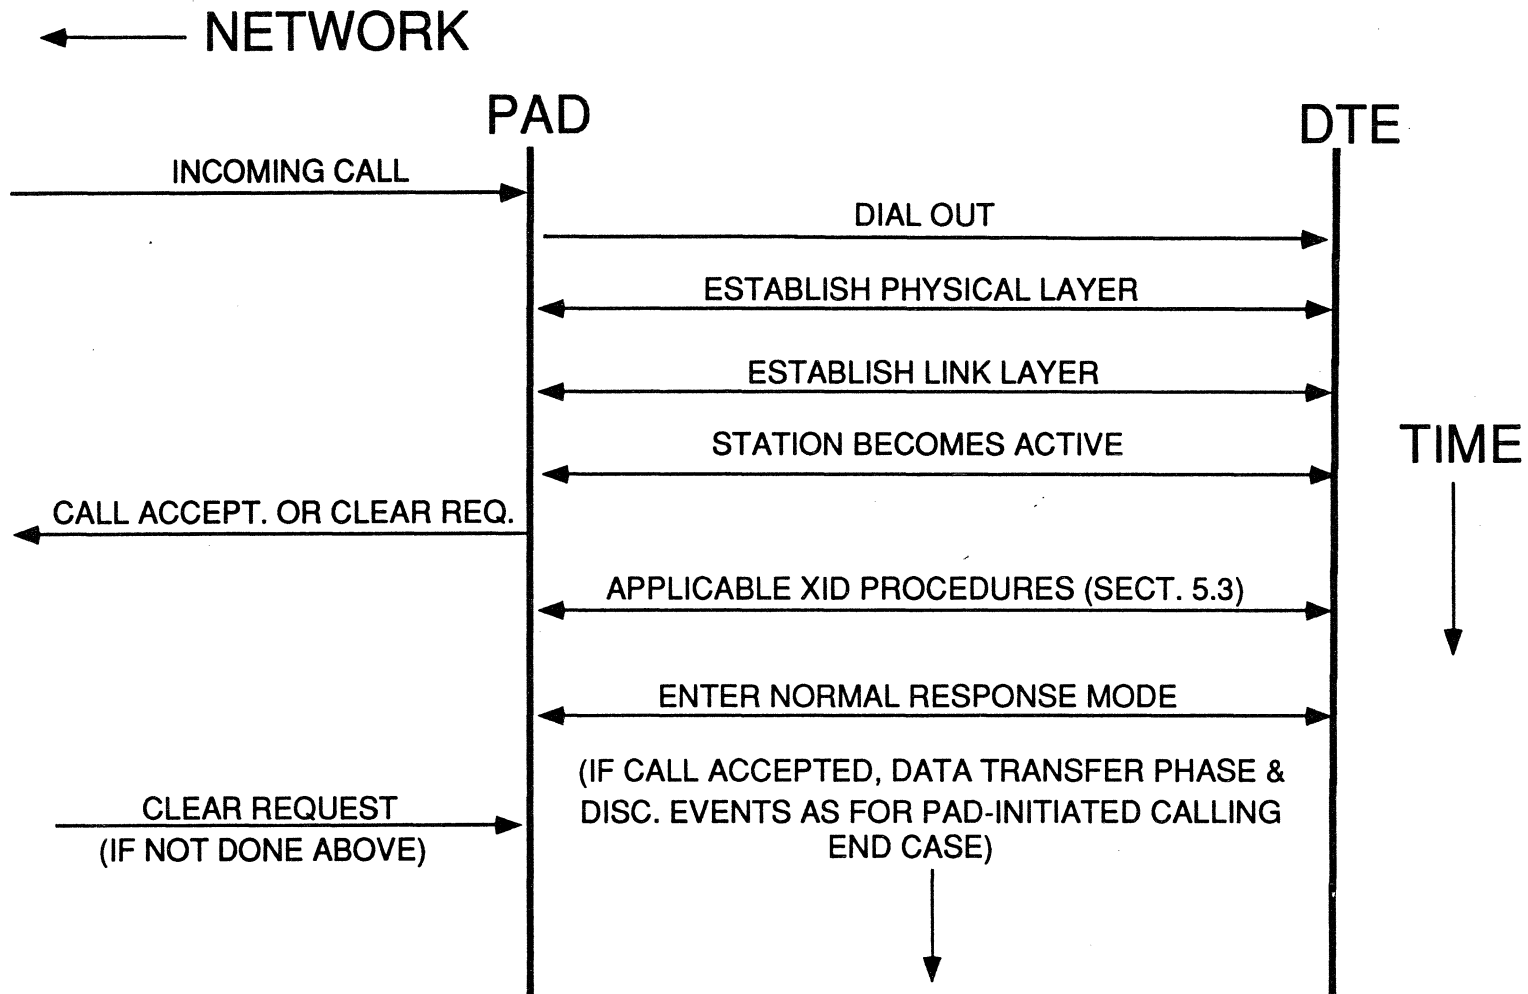


Figure 5-3. Time Relationships Among Link Layer and SDLC Establishment Events (Sheet 4 of 4)

Table 5-1. Valid SDLC PAD Interface Configuration Combinations

Local Station Type Served	PAD-Initiated SVC (Auto Call)		PVC	DTE-Initiated SVC	
	Dial-In	Dedicated		Dial-Out	Dedicated
Secondary	OK*	OK	OK	OK	OK
Primary	OK*	OK	OK	OK	OK
Negotiable	OK*	OK	OK	OK	OK

* Private dial-in ports.

Table 5-2. Two-Digit (Partial) VC Profiles With a Uniform Definition Across PPSNs

Prof. No.	Description	CUD Field in Hex	Rev. Chg?	Pkt Size	Wind. Size	Throughput Class**
10	ISO Diag., RC, Interactive, 128 octet	C3	Yes	128	2	2400
11	ISO Diag., RC, File Trans., 256 octet	C3	Yes	256	2	9600
12	ISO Diag., No RC, Interactive, 128 octet	C3	No	128	2	2400
13	ISO Diag., No RC, File Trans., 256 octet	C3	No	256	2	9600
14	SNA Diag., RC, Interactive, 128 octet	CB	Yes	128	2	2400
15	SNA Diag., RC, File Trans., 256 octet	CB	Yes	256	2	9600
16	SNA Diag., No RC, Interactive, 128 octet	CB	No	128	2	2400
17	SNA Diag., No RC, File Trans., 256 octet	CB	No	256	2	9600
18-29	[Reserved]					

RC =reverse charged

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

6. SDLC/VC MAPPINGS AND STATE CHANGES

This section addresses frame and packet mapping for the SDLC and X.25 packet layer protocols, respectively. It also covers queue handling and state changes.

6.1 Mapping SDLC Frames to/from Packets

The format and use of the supported QLLC qualified Data packets should conform to the IBM X.25 interface documents cited earlier (in particular, Section 4 of the general information manual^[8] and Section 8 and Appendix O of the architecture reference^[9]). For implementations of the mapping between frames and QLLC qualified Data packets on the VC connection to and from 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 packet's 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 (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 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 that 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 VC (cause code x'80' and diagnostic code decimal 161 or cause code x'00' and diagnostic code decimal 163).

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 the IBM X.25 interface architecture reference cited earlier^[9].

2. The PAD performs mapping between the SNRM(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. If supported by the PAD, mapping may also be performed for SIM, RIM, and UI frames, as specified in Section 3.2.2.1. 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 exceptions of packets containing user data as specified in item 1 and the special case of mapping RR frames to QRR packets (see item 4 of Section 3.2.2.1), all other packet types (such as RNR and REJ) are handled (and possibly discarded) within the VC service with no mapping to frames on the corresponding local SDLC interface. Tables 6-1 through 6-4 specify the relationships between the mappable frames and packets and summarize the actions taken by the PAD on receipt of a specific frame or QLLC packet type. Tables 6-1 and 6-3 summarize frame formats, frame generation triggers, and actions to be taken by the PAD on receipt of each frame. Tables 6-2 and 6-4 do the same for each QLLC qualified Data packet type. The handling of unqualified Data packets and packets

other than Data packets has already been addressed above.

3. If a PAD receives a frame or QLLC qualified Data packet that is not valid given the configured station role (indicated by "NA" in Tables 6-1 and 6-2), the PAD should adopt one of two courses of action in a consistent manner:
 - Map the frame to the corresponding QLLC packet, or vice versa, in the normal way and allow the receiving DTE to respond to the error on an end-to-end basis
 - Respond to the invalid frame or packet immediately (across the local SDLC interface or the VC, respectively) without mapping in a manner consistent with SDLC procedures (e.g., respond with a DM/QDM or FRMR/QFRMR; ignore the invalid frame/packet).

The PAD should not take any local action to prevent or respond to retransmitted commands (i.e., because of timeouts). The PAD should perform normal mapping and allow frame retransmission to be handled by the DTEs on an end-to-end basis. (It is assumed that the timeout values configured for the primary station correctly reflect applicable turnaround time for the virtual circuit so unnecessary retransmissions and duplicate responses do not normally occur.)

6.2 Send Queues and Flow Control: SDLC and Virtual Circuit

The following requirements pertain to send queues and flow control:

1. As discussed in Section 3.2.2.2, 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 VC 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 VC. When a complete, valid, in-sequence frame is received across the local SDLC interface, the PAD should map it to one (or more, if segmenting is required) packets, which are added to the corresponding VC queue.
2. VC send queues should be emptied in accordance with normal VC 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 3.2.2.2. 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 I-Frames on the local SDLC interface, independently, each in accordance with the procedures of the VC service and SDLC link interface, respectively. However, flow control conditions on the local SDLC interface and within the corresponding VC 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 VC 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 VC.

6.3 Relationships Between SDLC Interface and Corresponding VCs

Section 6.1 and Tables 6-1 through 6-4 specify the various mappings between SDLC frames and VC packets that are to be performed by the PAD. They also specify the mapping, frame generation, and status update actions to be taken by the PAD in response to various events and the receipt of specific SDLC frames and QLLC packets. This section addresses PAD responsibilities associated with exception conditions, state changes, and diagnostics with implications across both the local SDLC interface and the associated VC connections.

Specifically, this section covers requirements related to:

- Maintaining, establishing, and breaking the connection for local SDLC interfaces
- Conditions that result in the PAD resetting PVCs, clearing calls, and/or establishing calls
- PAD actions at the local SDLC interface that are triggered by resetting or clearing of the associated VCs.

Many of the following requirements reflect the need to alert SDLC stations (even if only indirectly) of any failure or status change at the remote end of a VC connection or in the VC connection itself. This is particularly important because the SNA entity associated with an SDLC station is typically "unaware" that a dedicated SDLC connection has been replaced by a VC connection mediated by PADs.

The following requirements apply:

1. The SDLC PAD should maintain the physical layer in the operational state, when possible, for dedicated SDLC interfaces.
2. The SDLC PAD should maintain the physical layer in the operational state, when possible, for dial-up SDLC interfaces, unless a condition that calls for breaking the dial connection exists (as specified below).
3. The SDLC PAD should maintain the link layer in the operational state when the physical layer is operational and, for dial-up interfaces, when the dial connection is established. ("Operational state" for the SDLC link layer does not imply that all or any of the stations on the link are necessarily in the normal response mode.)
4. For dial-out interfaces, the SDLC PAD establishes a dial connection:
 - a. To "complete" an SVC to a called address which corresponds to an E.164 PSTN or voiceband ISDN destination
 - b. Before sending a Call Request for a DTE-initiated SVC whose calling DTE address corresponds to an E.164 PSTN or voiceband ISDN originator.
5. The SDLC PAD breaks an established dial-up connection (and clears the associated SVC, if it is still established) if one of the following events occurs:
 - a. A configured XID check fails, for dial-in connections (see Section 5.3, item 8)
 - b. The associated VC is reset (with cause "out of order" or "network out of order") or is cleared
 - c. The carrier is disrupted or lost for more than 2 seconds (see Section 3.1).

6. For private, dial-in interfaces, the SDLC PAD should support the ability to initiate an SDLC XID exchange in response to establishment of the dial connection as a configuration option (see Section 5.3, item 4).
7. Initiating link layer mode transitions (e.g., from normal disconnected mode to normal response mode) is the responsibility of the primary station DTE (i.e., via appropriate QLLC packets across the VC). The SDLC PAD is not responsible for independently initiating such mode transitions.
8. The SDLC PAD should update the link layer mode status it maintains in memory for the appropriate (actual or emulated) secondary station on the SDLC interface in response to SDLC frames and QLLC packets received, as specified in the "PAD Actions" columns of Tables 6-3 and 6-4. The PAD's assumed mode for a secondary station should be reset to "disconnected mode" when it receives a DM/QDM, SNRM(E)/QSM, or UA/QUA (the last pair, only when in response to a previously transmitted DISC) for that station. The PAD's assumed mode for a secondary station should be reset to "normal response mode" when it receives a UA or QUA in response to a previously transmitted SNRM(E).
9. The PAD should clear an SVC and reset a PVC (with the appropriate cause and diagnostic codes) if a network problem or other condition occurs at the packet layer that requires clearing or resetting according to Recommendation X.25 for that VC (e.g., invalid packet type or packet too long).
10. The PAD should clear each established SVC affected with cause "out of order" (and not attempt to automatically reestablish the call) and should reset each PVC affected with cause "out of order" if any of the following conditions/events is detected:
 - a. The physical layer or link layer of the corresponding SDLC interface is declared non-operational (e.g., DTE power off)
 - b. A dial-up circuit connection in the path is disconnected
 - c. An error condition is detected on the local SDLC interface that is not recoverable at the link layer
 - d. The associated secondary station does not acknowledge polls after the established number of retries (the retry maximum should be BCC-settable at least to values in the range 0 to 5; see Section 3.2.2.2).
11. The SDLC PAD should reset the affected PVC with cause "remote DTE operational" if one of the following events occurs:
 - a. The physical layer or link layer of the corresponding SDLC interface is declared operational (e.g., DTE power on and successful SDLC operation) following a period of non-operational status or a link layer error condition that could not be recovered at the link layer
 - b. The associated secondary station responds validly to an "inactive station" poll after a period of inactivity (e.g., see Section 3.2.2.2).
12. When a PAD issues a Clear Request, Reset Request, or Restart Request acting on behalf of the network (i.e., as the DCE), it should use the diagnostic codes specified in Appendix E of the IBM X.25 interface architecture reference^[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 VC if it were present in place of the PAD and attached SDLC DTE), the DTE diagnostics in Appendix H of that IBM document should be used; part H.2 applies to connections over which QLLC is supported.

13. In addition to any other necessary condition, the SDLC PAD may initiate a Call Request only if the local SDLC interface is operational and any T_{d1} or T_{d2} timers that may have been activated (see later item in this section) have expired.
14. Subject also to the above requirement, the SDLC PAD should initiate a call request, as follows:
 - a. For SDLC interfaces/stations configured for PAD-initiated SVCs in response to a transition from link layer nonoperational to link layer operational
 - b. For SDLC interfaces/stations configured for DTE-initiated SVCs in response to a valid SVC establishment request received from an authorized user on an asynchronous control interface
 - c. For SDLC interfaces/stations configured for auto call service, as specified in Sections 5.4.2.1 and 5.4.2.2.
15. The SDLC PAD should modify its local polling and frame handling behavior for the affected SDLC secondary station address(es) during any periods of modified operation:
 - 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, or until expiration of the T_{d3} timer specified later, whichever is later
 - Following resetting of a PVC or SVC with any cause other than "out of order," "network out of order," "remote DTE operational," and "network operational," and until expiration of the T_{d3} timer specified later
 - Following clearing of an SVC until a new VC is established in association with the secondary station (in some cases, a T_{d1} or T_{d2} timer must expire before a new SVC may be established, as specified below).

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

During this period of modified operation, and until a new VC is established and associated with the affected secondary station (or the established PVC is again declared operational), the SDLC PAD:

- a. Should not transmit any frames with the affected secondary station address across the local SDLC interface, unless the transmission is directly associated with the process of eliminating the cause of modified operation (e.g., sending "inactive station" XID polls to a station that has been declared inactive because of failure to respond to polls)
- b. Should respond to all poll commands with a DM and should not acknowledge or respond to any other frames received over the local SDLC interface, unless the received frame is directly associated with the process of eliminating the cause of modified operation (e.g., receiving a response to an "inactive station" poll)
- c. Should discard any frames received and any (as yet unmapped) frames and (already mapped) packets queued for transmission over the affected VC.

Following the period of modified operation, the SDLC PAD should return to normal polling and frame handling operation.

In the case of a VC disruption, locally acknowledged frames can nevertheless be lost on an end-to-end basis. When an SVC is cleared, the next SVC to be established in association with the same (actual or emulated) secondary station may correspond to a different LU or DTE and undoubtedly will not correspond to the same point in the higher layer SNA session as existed just before the disruption took effect. Since there are no clearly defined mechanisms for signaling the occurrence of X.25 VC disruptions over affected SDLC interfaces, it is important for the SDLC PAD to indirectly alert SDLC DTEs attached to the PAD of these disruptions. For dial-up interfaces, the alerting is accomplished by disconnecting the dial-up connection, as specified earlier.

For dedicated SDLC interfaces, the alerting is accomplished by requiring a minimum duration for the period of modified operation, as specified in the following items:

- If the VC disruption was caused by the clearing of an SVC, the period of modified operation must last for a period of at least T_{d1} seconds if the clearing was initiated by the network or the remote end, and for the period of at least T_{d2} seconds if the clearing was initiated (or specified) by the local DTE.
- If the VC disruption consists of a reset, the period of modified operation must last for a period of at least T_{d3} seconds.

Each of these three timer values is a configuration option on dedicated SDLC interfaces that should be capable of being set to at least the values in the range of 0 to 120 seconds. By using the " T_d " timers, even if the primary station "thinks" it is directly connected to the remote end secondary station, it will receive a clear indication over the SDLC interface of the disruption in the intervening VC connection (i.e., no response to any commands sent for an extended period of time). (The appropriate values for these T_d timers depend on the timeout and retry parameter values configured in the FEP for the corresponding SDLC interface. These minimum-duration disruptions may also require manual intervention to re-enable establishment of sessions unless automatic restarts are configured at the host end.)

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)	QSM	QSM (NA)
DISC	QDISC	QDISC (NA)
XID	QXID (C)	QXID (R)
TEST	QTEST (C)	QTEST (R)
RR	QRR (C) [Note 3]	- (NA)
RD	QRD (NA)	QRD
UA	QUA (NA)	QUA
DM	QDM (NA)	QDM
FRMR	QFRMR (NA)	QFRMR
SIM*	QSIM	QSIM (NA)
RIM*	QRIM (NA)	QRIM
UI*	QUI (C)	QUI (R)

Legend & Notes

- (C) Command version (originated by primary station end)
- (R) Response version (originated by secondary station end)
- (NA) Not Applicable or Invalid. See item 3 of Section 6.1 for the correct treatment of such received frame (“-” indicates that “RR” received from the secondary station end is never mapped).
- * These mappings are not specified in QLLC; support is not required.

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 VC LCN.
2. Frame control field and information field (if present) conform to SDLC specifications; information field maps to/from corresponding QLLC packet fields.
3. A locally received RR is mapped to a QRR sent to the remote end only for a PAD serving a negotiable role station that sends the RR to the PAD immediately after the PAD has received a SNRM(E) from that station and subsequently delivered a UA to that station.

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) %	SNRM(E) (NA)
QDISC	FF	53	No	DISC	DISC (NA)
QXID (C)	FF	BF	Yes	XID	XID (NA)
QXID (R)	xx	BF	Yes	XID (NA)	XID
QTEST (C)	FF	F3	Yes	TEST	TEST (NA)
QTEST (R)	xx	F3	Yes	TEST (NA)	TEST
QRR (C)	FF	F1	No	RR***	RR (NA)
QRR (R)	xx	F1	No	RR (NA)	RR (NA)***
QRD	xx	53	No	RD (NA)	RD
QUA	xx	73	No	UA (NA)	UA
QDM	xx	1F	No	DM (NA)	DM
QFRMR	xx	97	Yes	FRMR (NA)	FRMR
QSIM %%	FF	17	No	SIM	SIM (NA)
QRIM %%	xx	17	No	RIM (NA)	RIM
QUI (C) %%	FF	13	Yes	UI	UI (NA)
QUI (R) %%	xx	13	Yes	UI (NA)	UI

Legend & Notes

- (C) Command version (originated by primary station end)
- (R) Response version (originated by secondary station end)
- (NA) Not Applicable or Invalid. See item 3 of Section 6.1 for the correct PAD treatment of such a received QLLC packet.

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.
- *** It is expected that the QRR will be sent only as a command from the primary end, and only for peer-to-peer connections for which the RR is used to explicitly signal that data transfer may commence.
- % Use of SNRM or SNRME depends on SDLC interface configuration.
- %% These packet types are not specified in QLLC; these extensions of QLLC are specified for multivendor compatibility if supported by the PAD.

Table 6-3. SDLC Frames: Summary of PAD Send Conditions & Actions on Receipt (Sheet 1 of 3)

Frame Type	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
SNRM(E)	Never	Mapped from incoming QSM	Map to outgoing QSM; update station status to disconnected mode	See Section 6.1, item 3
DISC	Never	Mapped from incoming QDISC	Map to QDISC	See Section 6.1, item 3
XID	See Section 5.3	See Section 5.3	See Section 5.3	See Section 5.3
TEST	Mapped from incoming QTEST response	Mapped from incoming QTEST command	Map to QTEST command	Map to QTEST response
RR	Frame window update poll response if not done by I-frame and no receive flow control	Frame window update (poll) if not done by I-frame and no receive flow control. Also, if receive QRR from "primary" peer end to signal data transfer may begin	Update frame window; end transmit flow control state if appropriate; if peer-to-peer connection, map to QRR if RR was in response to UA sent as acknowledgment of SNRM(E)	Update frame window; end transmit flow control state if appropriate
RNR	To signal receive flow control as part of response	To signal receive flow control. Also, if peer-to-peer connection, immediately after UA received in response to SNRM(E) sent and until QRR or data packet received from remote primary end	Update frame window; enter transmit flow control	Update frame window; enter transmit flow control

Table 6-3. SDLC Frames: Summary of PAD Send Conditions & Actions on Receipt (Sheet 2 of 3)

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 frame window, (re)transmit I frames as requested	Update frame window, (re)transmit I frames as requested
RD	Mapped from incoming QRD	Never	See Section 6.1, item 3	Map to outgoing QRD
UA	Mapped from incoming QUA	Never	See Section 6.1, item 3	Update station status, if in response to SNRM(E) or DISC; Map to outgoing QUA; If peer-to-peer connection and in response to SNRM(E) sent, send RNR to prevent local DTE from starting data transfer until remote primary acknowledges.

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

Frame Type	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
DM	Mapped from incoming QDM	Never	See Section 6.1, item 3	Update station mode; Also, map to outgoing QDM
FRMR	Mapped from incoming QFRMR	Never	See Section 6.1, item 3	Map to outgoing QFRMR
I	Mapped from unqualified Data packet, in response to poll	Mapped unqualified Data packet, according to poll procedures	Map to unqualified Data packet (sequence) over VC	Map to unqualified Data packet (sequence) over VC
Other	Never	Never	Respond as to invalid frame; see sect. 6.1, item 3	Respond as to invalid frame; see sect. 6.1, item 3
SIM *	Never	Mapped from incoming QSIM	Mapped to outgoing QSIM	See Section 6.1, item 3
RIM*	Mapped from incoming QIRM	Never	See Section 6.1, item 3	Map to outgoing QIRM
QUI*	Mapped from incoming UI response	Mapped from incoming UI command	Map to outgoing QUI command	Map to outgoing QUI response

* Support of these frames and associated mapping is not required. See item 3 of Section 3.2.2.1 for nonsupport treatment alternatives.

Table 6-4. QLLC Packets: Summary of PAD Send Conditions & Actions on Receipt (Sheet 1 of 2)

QLLC Packet	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
QSM	Mapped from incoming SNRM(E)	Never	See Section 6.1, item 3	Map to SNRM(E); update station to disconnected mode
QDISC	Mapped from incoming DISC	Never	See Section 6.1, item 3	Map to DISC
QXID	See Section 5.3	See Section 5.3	See Section 5.3	See Section 5.3
QTEST	Mapped from incoming TEST command	Mapped from incoming TEST response	Map to TEST response	Map to TEST command
QRR	Only if peer-to-peer connection and RR received from "primary" in direct response to UA sent as acknowledgment to SNRM(E)	Never	Update window and flow control receive accordingly	Update window and flow control receive accordingly; if peer-to-peer connection, send RR to local DTE to signal that data transfer may start
QRD	Never	Mapped from incoming RD	Map to RD	See Section 6.1, item 3
QUA	Never	Mapped from incoming UA	Map to UA and update station mode based on frame responding to	See Section 6.1, item 3

Table 6-4. QLLC Packets: Summary of PAD Send Conditions & Actions on Receipt (Sheet 2 of 2)

QLLC Packet	When PAD Transmits		PAD Action(s) if Received	
	At Primary End	At Secondary End	At Primary End	At Secondary End
QDM	Never	Mapped from incoming DM	Map to DM and update station status to disconnected mode	See Section 6.1, item 3
QFRMR	Never	Mapped incoming FRMR; Also, if invalid/unsupported QLLC or QLLC-like packet received (see Section 6.1, item 3)	Map to FRMR	See Section 6.1, item 3
UDP (Data)	Mapped from incoming I-Frame; update window	Mapped from incoming I-Frame; update window	Map to I-Frame; update window	Map to I-Frame; update window and flow control status
QSIM*	Mapped from incoming SIM	Never	See Section 6.1, item 3	Map to SIM
QRIM*	Never	Mapped form incoming RIM	Map to RIM	See Section 6.1, item 3
QUI*	Mapped from incoming UI command	Mapped from incoming UI response	Map to UI response	Map to UI command

* Support of these qualified packets and associated mapping not required. See item 3 of Section 3.2.2.1 for nonsupport treatment alternatives.

7. GLOSSARY

7.1 Definition of Terms

Asynchronous Balanced Mode. 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). For SNA equipment, typically used only on token ring connections. (The LAPB link layer protocol of X.25 uses this type of balanced operation.)

Authorization 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).

Binary Synchronous Communications. A link layer 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.

Boundary Network Node. A (type 4 or 5) node that uses network addresses for routing and provides support to attached peripheral nodes.

Call User Data. 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.

Cluster Controller. A device that 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*.

Communication Controller. See entry for *front end processor*.

Display System Protocol. 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.

Enhanced Logical Link Control. 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.

Front End Processor. A general term for what is called a *communication controller* within the IBM SNA framework. This is a (type 4) node that 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.

High-level Data Link Control. An ISO defined link layer protocol that includes SDLC and X.25 LAPB as subsets or specific classes of procedure.

Initialization Mode. See description under *normal response mode*.

Logical Link Control. 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.

Logical Unit. 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.

NCP Packet Switching Interface. See entry for *X.25 NCP Packet Switching Interface*.

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.

Network Addressable Unit. A source or destination of information being passed through an SNA network. There are three types of NAU: LU, PU, and SSCP.

Network Control Program. An IBM program that runs in an FEP and provides communication controller support.

Network Interface Adapter. An IBM protocol conversion device that 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.

Normal Disconnected Mode. See description under *normal response mode*.

Normal Response Mode. 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.

Physical Services Header. A protocol used by IBM Network Interface Adapters to encapsulate SNA communications within X.25 packets so they can be transported across an X.25 packet network. PSH corresponds to one of the LLCs defined by IBM.

Physical Unit. 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.

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.

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).

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.

Qualified Logical Link Control. 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.

Secondary. See the description of *primary*.

Synchronous Data Link Control. The SNA link layer 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.

System Services Control Point. 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.

Virtual Circuit 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.

Virtual Telecommunication Access Method. 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 that is expected to have full IBM support into the future.

X.25 NCP Packet Switching Interface. IBM software resident in an FEP that permits that FEP to directly support X.25 interfaces.

7.2 Acronyms

ABM	Asynchronous Balanced Mode
ABME	Asynchronous Balanced Mode Extended
AC	Access Concentrator
ACF	Advanced Communication Function

AFM	Advanced Function Modem
ARF	Activation Response-Lag Factor
BCC	Bellcore Client Company
BNN	Boundary Network Node
BSC	Binary Synchronous Communication (Bisynchronous)
CCITT	International Telegraph and Telephone Consultative Committee
CD	Data Terminal Ready (EIA 2320 interchange circuit)
CPE	Customer Premises Equipment
CUD	Call User Data
CUG	Closed User Group
DISC	Disconnect
DM	Disconnected Mode
DSP	Display System Protocol
ELLC	Enhanced Logical Link Control
FEP	Front End Processor
FRMR	Frame Reject
HDLC	High-level Data Link Control
HPAD	Host PAD
IC	Interexchange Carrier
LAPB	Link Access Procedure - Balanced
LCN	Logical Channel Number
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
OA&M	Operations, Administration, and Maintenance
OSI	Open System Interconnection
OTGR	Operations Technology Generic Requirements
PAD	Packet Assembler/Disassembler

PPSN	Public Packet Switched Network
PPSNGR	PPSN Generic Requirements
PS	Packet Switch
PSH	Physical Service Header
PSPDN	Packet Switched Public Data Network
PSTN	Public Switched Telephone Network
PU	Physical Unit
PVC	Permanent Virtual Circuit
QLLC	Qualified Logical Link Control
RD	Request Disconnect
RNR	Receiver Not Ready
SDL	System Description Language
SDLC	Synchronous Data Link Control
SNA	Systems Network Architecture
SNRM	Set Normal Response Mode
SPAD	Symmetric PAD
SSCP	System Services Control Point
SVC	Switched Virtual Circuit
TA	Technical Advisory
TPAD	Terminal PAD
TR	Technical Reference
UA	Unnumbered Acknowledgment
UI	Unnumbered Information
VC	Virtual Circuit
VTAM	Virtual Telecommunications Access Method
XID	Exchange Station Identification

8. REFERENCES

1. *Public Packet Switched Network Generic Requirements (PPSNGR)*, TR-TSY-000301, Bellcore, Issue 2, December 1988.
2. *Systems Network Architecture Concepts and Products*, IBM Document GC30-3072-3, October 1986.
3. *Systems Network Architecture Technical Overview*, IBM Document GC30-3073-2, September 1986.
4. *Operations Technology Generic Requirements (OTGR)*, TR-TSY-000439, Issue 2, Bellcore, February 1988; plus Revision 1, November 1988.
5. *Synchronous Data Link Control Concepts*, IBM Document GA27-3093-3, June 1986.
6. *Systems Network Architecture Format and Protocol Reference Manual: Architecture Logic*, IBM Document SC30-3112-2, November 1980.
7. *Format and Protocol Reference Manual: Architecture Logic for Type 2.1 Nodes*, IBM Document SC30-3422-0, December 1986.
8. *The X.25 1984 Interface for Attaching SNA Nodes to Packet-Switched Data Networks General Information Manual*, IBM Document GA27-3761-0, November 1986.
9. *The X.25 1984 Interface for Attaching SNA Nodes to Packet-Switched Data Networks Architecture Reference*, IBM Document SC30-3409-0, December 1987.
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 this 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
P.O. 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
1 (800) 521-CORE
(201) 699-5800 (for foreign calls)

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

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 perform protocol conversion between an X.25 type Virtual Circuit (VC) 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 a relatively cost-effective PAD implementation and network transparency for such functions as customer network management. However, certain services cannot be provided within the PPSN if only the SDLC layer is processed:

- Supporting separate VC 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 as to what is 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, separating 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 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 an SNA DTE 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 VC 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 remainder of this Appendix 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 when its services are required. Once this VC 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 VCs established and SNA services provided.

Figure A-1 suggests the outline 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 VCs. 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 VC 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-to-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.




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

Appendix B

Prototype Cases Used for Performance Specifications*

Assumptions and Conditions Applying to All Cases

1. All SDLC interfaces on the PAD are configured to support secondary stations. Three quarters of the interfaces are configured for multipoint operation. The remaining one quarter are configured for point-to-point operation.
2. One half of the SDLC interfaces are configured for the maximum number of secondary stations that the interface can support at the given line speed. One quarter of the SDLC interfaces are configured for one half the maximum number of secondary stations that the interface can support at the given line speed. The remaining one quarter of the SDLC interfaces are configured for one secondary station. Nonintegral fraction remainders are consolidated and the resulting integral interface is assigned to the contributing category with the largest number of stations per interface.
3. All configured secondary stations are operational and are already in the NRM mode. All secondary stations respond to polls immediately at the given line speed (e.g., propagation delay and DTE frame handling delays may be ignored). However, none of the secondary stations have any data (I frames) to transmit and there are no I frames queued for transmission to any of the secondary stations.

Case 1: Low Line Speeds

1. All SDLC interfaces operate at 1200 bps.

Case 2: Moderate Line Speeds

1. All SDLC interfaces operate at 4800 bps.

Case 3: High Line Speeds

1. All SDLC interfaces operate at 56000 bps.

* See item 2 of Section 3.2.2.2 for relevant performance specifications.

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
- 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 (nonliteral) entries or fields are enclosed in angle brackets (< description >)
- Options are enclosed in square brackets ([options])
- Commas separate options within square brackets that 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 parentheses) 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 appendix.

A request may be entered on the same line as the escape character, as illustrated below (assuming the system is 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 a 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> <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 and destination address fields associated with a normal asynchronous call request (see PPSNGR, Section 3.3). (Note that the normal DTE-supplied user data field is not included here since such DTE-to-DTE data would present frame sequencing problems for mapping by the PAD at the remote end SDLC interface.) 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.

If VC profiles are supported (see Section 5.5 for description and format), the destination address field can be used to specify the VC profile number. If the abbreviated addressing format is used (see PPSNGR, Section 3.3.3.11.13), the first character of the address field is a period (“.”). If the character following the period is the number sign (“#”), one or two decimal digits should follow that specify the VC profile number. 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.

Whether an alphanumeric extension of the VC profile number concept should be adopted is under consideration (e.g., “.&<alphanumeric string>” would represent a VC profile while any string starting with other than a “#” or “&” following the “.” would represent a simple abbreviated address). Also under consideration (and beyond the scope of this document) is whether the ability to designate a VC profile in the destination address field should be generalized to apply to non-SDLC (e.g., X.28) interfaces as well.

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>,A<orig. addr.>[,X<hex CUD>]

There are two required and one 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 dedicated interfaces, the authorization code is an alphanumeric string of 6 to 10 characters that must be unique within the network element (i.e., AC or PS). For dial-up interfaces, the code should be 10 characters long and unique within the PPSN. In both cases, the code must contain at least one decimal digit or special symbol, 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> or comma that terminates the code value.)

For each SVC establishment request, the submitted authorization code is checked against the code(s) authorized for the SDLC interface or SDLC interface/station combination identified as the calling DTE for the SVC. The request is processed only if there is an authorization code match. In addition, for dial-up control interfaces, the authorization code first submitted at dial-in is checked against the list of codes configured for the dial-up port. If it is a valid code, the PSTN address associated with that code is used to establish a dial-out connection control interface for added security, before any control mode request is honored. Each dial-up asynchronous interface supporting the control function should be configurable with at least 100 valid authorization code values and associated PSTN dial-back addresses.

2. 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 PPSNGR, Section 3.3). However, this address is constrained to be a PPSN X.121 address of an SDLC interface on the same AC or PS.
3. Each SVC establishment request (either explicitly, or via a VC profile designation) includes a specification of the first 1 to 16 octets of the call user data (CUD) field, expressed as twice that number of hexadecimal digits (semioctets), following the code letter "X". Note that user data specified as alphanumeric by the DTE (i.e., following a "D" or "P" at the end of the normal portion of the X.28 call setup command) is not permitted because of problems in mapping such data to the SDLC interface. The hexadecimal CUD octets are used to specify the protocol identifier and secondary station PU/LU characteristics, as specified in item 4 of Section 5.4.2.1.

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

```
SDLCntl> Z5551234JM$,A5551111,XC3,T3299,R-132505559999<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 and the CUD field should contain 4 octets with the hexadecimal value C3.

SDLC PAD SVC Clear Request Format

While in the SDLC PAD control function mode, the user may request that any existing SVC for which the user has control authorization be cleared. The format of the clear request is as follows:

```
CLR|clr:<address of calling DTE>,<logical channel number><CR>
```

where the originating DTE address and logical channel number uniquely identify the SVC from among those that are currently active and those 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). Clear Request packets generated as a result of such commands have the cause "DTE originated" and diagnostic code decimal 0. The cause code value is x'00' if the first octet of the CUD of the Call Request packet had been x'C3' and is x'80' if the CUD first octet had been x'CB'.

Message Display Request Formats

The default procedure displays 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[:<address of calling DTE>,<logical channel number>]<CR>

If a specific VC is not identified via a DTE address and logical channel number, the above command applies to all VCs for "off" and all VCs established or explicitly requested for display during this session for "on."

Appendix D

Control Interface Display Formats for Status and Clearing Messages

The control interface that 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 (noncontrol) 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 (T_d) BUFFER PERIOD (<time remaining>)

Appendix E

Recommended Formats for Menu-Oriented Command and 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
 - "Control a" moves the cursor to the beginning of the field
 - "Control e" moves the cursor to the end of the field
 - "Control k" erases the current field
 - Backspace ("Control h") moves the cursor to the previous field character position
 - Overwriting 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.

Figures E-1 through E-4 present the recommended formats for four display screens.

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 dedicated 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./

Figure E-1. Screen #0 Display Format

SDLC PAD CONTROL MODE

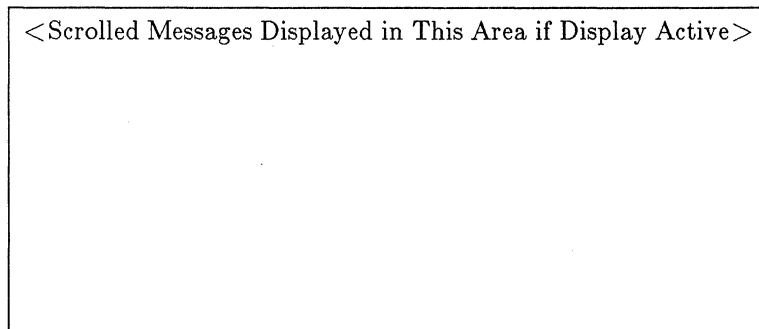
SCREEN #1: REQUEST/COMMAND SUMMARY

- 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) /

Figure E-2. Screen #1 Display Format

SDLC PAD CONTROL MODE

SCREEN #2: SVC ESTABLISHMENT

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>{min(9600,line speed)} D->O: <75,...,48000>{min(9600,line speed)}

Des. max. transit delay (ms): <1 to (2**16 - 1)>

CUD field: <even # of hex. digits>



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; fast select is not supported/

Figure E-3. Screen #2 Display Format

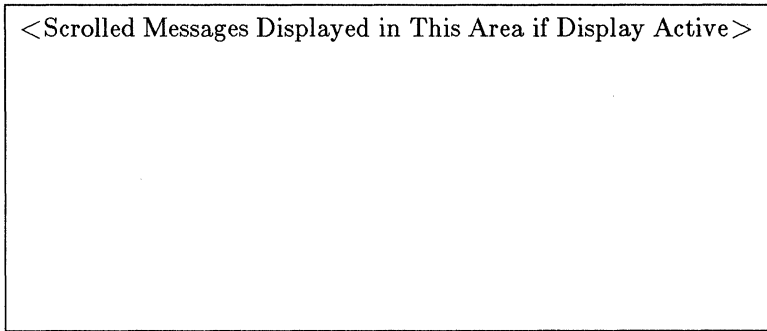
SDLC PAD CONTROL MODE

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}*

<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>

Figure E-4. Screen #3 Display Format

BELLCORE

PPSN Support of SNA/SDLC Interfaces
JSF/RAF

TR-TSY-000885
ISSUE 1, December 1989

READER RESPONSE FORM

You can help us make our document more useful. After you've used the document and have become familiar with it, please take a few minutes to fill out and return this form.

1. How did you use this document?

- To learn the procedures
- As a reference
- To find specific information for use right away
- To note updates and changes

Other: _____

2. How would you rate this document for organization of material?

- Excellent
- Very Good
- Good
- Fair
- Poor

Explain: _____

3. How would you rate this document in answering all your questions?

- Excellent
- Very Good
- Good
- Fair
- Poor

4. How would you rate this document for accuracy of information?

- Excellent
- Very Good
- Good
- Fair
- Poor

Explain: _____

5. Overall, how would you rate the usefulness of this document in meeting your information needs?

- Excellent
- Very Good
- Good
- Fair
- Poor

Identify any information that you feel should be included or removed.

Describe _____

Your job title: _____

Name (optional): _____

Company: _____

Address: _____

City and State: _____

Thank you for your cooperation. No stamp is necessary if mailed in the U.S.A.

NOTE: STAPLES CAN CAUSE PROBLEMS WITH AUTOMATIC MAIL-SORTING EQUIPMENT. PLEASE USE GUMMED TAPE TO SEAL THIS FORM.

Fold And Tape

Please Do Not Staple


Fold And Tape

ASD-492 (4-89)

From _____



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 9 PISCATAWAY, NJ

POSTAGE WILL BE PAID BY ADDRESSEE

BELLCORE

ATTN: TECHNICAL PUBLICATIONS QA COORDINATOR, PYA-1A283
P.O. BOX 1354
PISCATAWAY NJ 08855-9963



Fold And Tape

Please Do Not Staple

Fold And Tape