

SNASW DESIGN AND IMPLEMENTATION GUIDE

VOLUME 6 IN THE CISCO INTERNETWORKING DESIGN GUIDE SERIES

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
<http://www-europe.cisco.com>
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

**Americas
Headquarters**
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
<http://www.cisco.com>
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

**Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the
Cisco Connection Online Web site at <http://www.cisco.com>.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark
England • France • Germany • Greece • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Russia • Saudi Arabia • Scotland • Singapore
South Africa • Spain • Sweden • Switzerland • Taiwan, ROC • Thailand • Turkey • United Arab Emirates • United States • Venezuela

	About This Design Guide.	ix
	Intended Audience.	ix
	Customers with Migration Requirements	ix
	Customers with Enhanced Data Center Requirements.	x
	Credits.	x
	Obtaining Documentation.	x
	World Wide Web	x
	Documentation CD-ROM	x
	Ordering Documentation	xi
	Documentation Feedback	xi
	Obtaining Technical Assistance	xi
	Cisco.com.	xi
	Technical Assistance Center	xi
Chapter 1	Introduction	1-1
	What Is SNA Switching Services?	1-1
	SNASw Technology.	1-1
	Branch Extender.	1-2
	Enterprise Extender (HPR/IP).	1-3
	Dependent Logical Unit Requester/Dependent Logical Unit Server.	1-3
	SNASw Benefits.	1-4
	Scalability.	1-4
	Transport	1-4
	Usability, Serviceability, and Management	1-5
Chapter 2	Positioning SNASw and DLSw+	2-1
	Overview.	2-1
	DLSw+ Features	2-2
	Architecture	2-2
	Availability.	2-3
	Scalability.	2-3
	QoS	2-3
	Flexibility	2-3
	Management	2-4
	Cost	2-4
	SNASw Features	2-5
	Architecture	2-5
	Availability.	2-5
	Scalability.	2-5

	QoS.....	2-6
	Flexibility	2-6
	Management	2-7
	Cost.....	2-7
	Decision Criteria.....	2-7
	Is SNA Routing Required?	2-7
	Are You Ready?	2-7
	Network Design	2-8
	Case Study 1—No HPR over IP Transport (DLSw+ Only)	2-8
	Case Study 2—DLSw+ and SNASw.....	2-9
	Case Study 3—SNASw Only	2-10
	Conclusion	2-11
Chapter 3	Implementation of SNASw	3-1
	Getting Started	3-1
	When Is APPN Required?.....	3-1
	SNASw General Considerations	3-2
	Required SNASw Configuration.....	3-2
	Supported SNASw Link Types and Transport Options.....	3-6
	Native LAN Transport	3-6
	Virtual Token Ring Transport	3-7
	VDLC Transport.....	3-8
	Native IP DLC Transport	3-8
	SNASw Connection Network Support	3-9
	SNASw Connection Network Support for Branch-to-Branch Traffic.....	3-10
	IBM CS/390 Connection Network Support.....	3-11
	Enabling SNASw DLUR Support	3-11
	Customer Scenarios for SNASw Deployment	3-13
	Customers Migrating from APPN NN to SNASw.....	3-13
	Customers Migrating from Subarea SNA (FEP) to SNASw	3-22
	Customers Migrating from Token Ring to High-Speed Ethernet Campus LAN	3-24
Chapter 4	SNASw Migration Scenarios	4-1
	Scenario 1—Large Branch Network Migration to SNASw EE	4-1
	Business Requirements	4-1
	Network Analysis	4-1
	Design Rationale.....	4-2
	The Migration.....	4-3
	The Old Network versus the New Network	4-7
	Scenario 2—Data Center and Remote Branch Migration from APPN NN and DLSw to SNASw EE.....	4-8
	Business Requirements	4-8
	Network Analysis	4-8
	Design Rationale.....	4-9

	The Migration	4-9
	The Old Network versus the New Network	4-10
	Scenario 3—Data Center and Remote Branch Migration from IBM 950 NNs and APPN (PSNA) to SNASw	4-10
	Business Requirements	4-11
	Network Analysis	4-11
	Design Rationale	4-11
	The Migration	4-12
	The Old Network versus the New Network	4-13
Appendix A	Glossary	A-1
Appendix B	APPN Components and Features	B-1
	APPN Technology Overview	B-1
	APPN Components and Node Types	B-1
	NN	B-2
	NN Server	B-2
	EN	B-2
	LEN Node	B-3
	CNN	B-3
	CP	B-3
	BrNN	B-3
	HPR Node	B-3
	ICN	B-4
	DLUR/ DLUS	B-4
	Dependent and Independent LUs	B-4
	Border Node	B-5
	Migration Data Host	B-5
	Transmission Group	B-5
	VRN	B-5
	CDS	B-5
Appendix C	Enterprise Extender (HPR/IP) Sample Configuration.	C-1
	Overview	C-1
	SNASw Configurations	C-2
	SNASw Router	C-2
	DSPU Router	C-3
	CIP Router	C-4
	Host Definitions	C-5
	SNASw Verification Commands	C-7
Appendix D	SNASw Hardware and Software Requirements	D-1
	Supported Hardware	D-1
	Supported Software	D-1
	APARs	D-1
	Supported MIBs, RFCs, and Standards	D-2

Appendix E	Frequently Asked Questions about APPN-to-SNASw Migration	E-1
	General	E-1
	Migration Considerations	E-2
Appendix F	SNASw Performance	F-1
	Overview	F-1
	The Test Environment	F-2
	Results	F-3
	Branch Router HPR/IP Performance	F-3
	Data Center Routers Running HPR/IP with DLSw+	F-4
	Data Center Routers Running HPR/IP without DLSw+	F-4
	Comparison of SNASw Modes for Data Center Routers Running with DLSw+	F-5
	Comparison of SNASw and PSNA for Branch Router	F-6
	Comparison of SNASw and PSNA for RSP4	F-7
	Version Information	F-8
Appendix G	References and Recommended Reading	G-1

Figure 1-1	Branch Extender Solution.	1-2
Figure 1-2	End-to-End EE Solution	1-3
Figure 1-3	DLUR/DLUS Support.	1-4
Figure 2-1	Comparison of Availability Characteristics of DLSw+ and SNASw	2-2
Figure 2-2	DLSw+ Architecture	2-3
Figure 2-3	SNASw Architecture	2-5
Figure 2-4	DLSw+ Design	2-9
Figure 2-5	Combined SNASw and DLSw+ Design	2-10
Figure 2-6	SNASw Design	2-11
Figure 3-1	Required Configuration	3-3
Figure 3-2	SNASw and Connection Network	3-4
Figure 3-3	SNASw Uplinks and Downlinks	3-5
Figure 3-4	Native LAN Transport	3-7
Figure 3-5	Virtual Token Ring Transport.	3-7
Figure 3-6	VDLC Transport	3-8
Figure 3-7	Native IP DLC Transport.	3-9
Figure 3-8	VRNs and Connection Networks	3-10
Figure 3-9	Connection Network Branch-to-Branch Example.	3-11
Figure 3-10	SNASw DLUR Support	3-12
Figure 3-11	ARB Flow Control.	3-14
Figure 3-12	BX Network Design for Migration	3-17
Figure 3-13	EE Migration Model 1: DLSw+ to the Branch	3-19
Figure 3-14	EE Migration Model 2: EE to the Branch	3-21
Figure 4-1	Install Data Center WAN Routers.	4-3
Figure 4-2	Install Regional Backbone Router.	4-4
Figure 4-3	Install Regional Distribution Routers	4-5
Figure 4-4	Convert the Branches and then the Regions.	4-6
Figure 4-5	Decommission the Regional IBM 2216 Routers	4-7



Figure 4-6	Implement the EE Connection Network (VRN)	4-10
Figure 4-7	Migrating Regional Banks to BX/DLUR	4-12
Figure C-1	Network Topology	C-1
Figure F-1	Traffic Generation Setup	F-2
Figure F-2	Branch Routers Running HPR/IP	F-3
Figure F-3	Data Center Routers Running HPR/IP with DLSw+	F-4
Figure F-4	Data Center Routers Running HPR/IP without DLSw+	F-5
Figure F-5	Comparison of SNASw Modes for Data Center Routers Running with DLSw+	F-6
Figure F-6	Comparison of SNASw and PSNA in a Cisco 4700 Series Branch Router	F-7
Figure F-7	Comparison of SNASw and PSNA on the RSP4	F-8

About This Design Guide

Intended Audience

This design and implementation guide is for customers who want to learn more about the Cisco Systems SNA Switching Services (SNASw) solutions. This document includes design guidelines and sample configurations appropriate for all audiences. It assumes familiarity with networking and Cisco routers but does not assume mastery of either. In some examples, key configuration commands are given for better understanding or to clarify a point. However, this design guide does not contain either exact or complete configuration information. For this type of information, consult Cisco's primary, real-time support system (see the Obtaining Documentation and Obtaining Technical Assistance sections in this chapter).

The general business requirements listed in this chapter are your navigation aids. Use these aids to assist you in distilling from this design guide the essential information for evaluating any proposed plans to upgrade your network components or features. Each of the general business requirements is further supported with detailed information within the chapter *Implementation of SNASw*.

Customers with Migration Requirements

Customers with migration requirements fall into three major categories:

- Existing SNA customers who:
 - Have versions of Cisco IOS Software® with Advanced Peer-to-Peer Networking (APPN) Network Node (NN) who are facing end-of-engineering support
 - Have a need for IP support in their current corporate, SNA-based network
- Customers who currently depend on IBM technology (for example, routers, front-end processors [FEPs], and Token Ring switches) and wish to:
 - Reduce network cost through FEP replacement
 - Migrate from a Token Ring to an Ethernet network
- Customers who have large numbers of APPN NNs and are experiencing:
 - Scalability issues due to large amounts of broadcast traffic and topology updates in their networks
 - Serviceability issues in which it is difficult to track down the solutions
 - Usability issues in which it is becoming more difficult to configure the network

Customers with Enhanced Data Center Requirements

Customers with enhanced data center requirement fall into two major categories:

- Customers who want to establish high-speed Internet Protocol (IP) data centers by using:
 - Cisco 7500 Series routers with Channel Interface Processor (CIP) attachment to the mainframe
 - Cisco 7200 Series routers with Channel Port Adapter (CPA) attachment to the mainframe
 - Catalyst® 6500 Series switches with IBM Open Systems Adapter-Express (OSA-Express) attachment to the mainframe
- Customers who want to increase their network availability and require:
 - Parallel Sysplex for SNA application availability on host logical partitions (LPARs)
 - APPN High Performance Routing (HPR) for nondisruptive SNA session rerouting around link failures
 - Multi-Node Persistent Sessions (MNPS) for rapid recovery from host application and system failures by eliminating re-establishment of sessions and enabling transparent recovery by the application
 - IBM WebSphere

Credits

Many people have contributed to the creation of the design guide. Steve Koretsky, who is the technical owner, wishes to thank the following people for their input and support:

- Lisa Bobbitt
- Jan Buskirk
- Dan McCullough
- Sudhir Nath
- Terry Otto
- Ray Romney
- Rick Williams

Obtaining Documentation

The following sections provide sources for obtaining documentation from Cisco.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at the following sites:

- www.cisco.com
- www-china.cisco.com
- www-europe.cisco.com

Documentation CD-ROM

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or as an annual subscription.



Ordering Documentation

Cisco documentation is available in the following ways:

- Registered Cisco direct customers can order Cisco product documentation from the Networking Products MarketPlace at www.cisco.com/cgi-bin/order/order_root.pl.
- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store at www.cisco.com/go/subscription.
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco corporate headquarters (California, USA) at 408 526-7208 or, in North America, by calling 800 553-NETS (6387).

Documentation Feedback

If you are reading Cisco product documentation on the World Wide Web, you can submit technical comments electronically. Click Feedback in the toolbar and select Documentation. After you complete the form, click Submit to send it to Cisco. You can e-mail your comments to bug-doc@cisco.com. To submit your comments by mail, use the response card behind the front cover of your document, or write to the following address:

Attn. Document Resource Connection
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain documentation, troubleshooting tips, and sample configurations from online tools. For Cisco.com registered users, additional troubleshooting tools are available from the Cisco Technical Assistance (TAC) Web site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information and resources at anytime, from anywhere in the world. This highly integrated Internet application is a powerful, easy-to-use tool for doing business with Cisco. Cisco.com provides a broad range of features and services to help customers and partners streamline business processes and improve productivity. Through Cisco.com, you can find information about Cisco and our networking solutions, services, and programs. In addition, you can resolve technical issues with online technical support, download and test software packages, and order Cisco learning materials and merchandise. Valuable online skill assessment, training, and certification programs are also available. Customers and partners can self-register on Cisco.com to obtain additional personalized information and services. Registered users can order products, check on the status of an order, access technical support, and view benefits specific to their relationships with Cisco. To access Cisco.com, go to www.cisco.com.

Technical Assistance Center

The Cisco TAC Web site is available to all customers who need technical assistance with a Cisco product or technology that is under warranty or covered by a maintenance contract.

Introduction

What Is SNA Switching Services?

The corporate intranet is replacing the Systems Network Architecture (SNA) WAN and the IP-based enterprise network is evolving to a converged network carrying voice, video, and integrated data. SNA Switching Services (SNASw) is a new release of the Advanced Peer-to-Peer Networking (APPN) feature of Cisco IOS® Software that enables operators of enterprise networks to develop their IP infrastructures while meeting SNA routing requirements. Companies can continue to enhance and deploy a converged network based on IP while handling traffic destined for SNA-based applications without change to the application itself. SNASw provides SNA routing and, optionally, SNA transport over IP. Further, SNASw was designed to be easy to use and configure, and it includes several innovative diagnostic features.

SNASw, available in Cisco IOS Release 12.1 and later, replaces the function provided by the APPN network node (NN) feature of Cisco IOS Software. Cisco is discontinuing the APPN NN feature in Cisco IOS Software beginning with Release 12.1 because of its architectural complexity, scaling limitations, and manageability issues. Older versions of Cisco IOS Software prior to Release 12.1 are reaching their end of engineering (EOE) support as follows:

- Release 11.2—April 16, 2001
- Release 12.0—After March 2002


SNASw Technology

SNASw adheres to APPN architecture while offering integration with other Cisco IOS features. Networks may be designed with SNASw at the remote branch, data center, or distribution layers in conjunction with Data-Link Switching Plus (DLSw+). Traffic flowing upstream from SNASw to the data center can utilize any SNA transport facility such as Cisco Channel Interface Processors (CIPs), Channel Port Adapters (CPAs), or the IBM Open Systems Adapter-Express (OSA-Express). SNASw can help to eliminate the need for SNA routing by front-end processors (FEPs) in the network.

Branch Extender

SNASw implements the SNA node type Branch Network Node (BrNN).¹ A BrNN appears as an end node (EN) to an upstream NN—usually IBM Communications Server for S/390 (CS/390)—while providing NN services for downstream ENs and low-entry networking (LEN) nodes (see Figure 1-1). This support for BrNN is also commonly referred to as Branch Extender (BX).

Figure 1-1 Branch Extender Solution



BX enhances scalability and reliability of SNA routing nodes by greatly reducing topology updates and eliminating broadcast directory storms that can cause network instability. Customers that attempted to build large SNA routed networks found that when the number of NNs exceeded about 75, scalability problems became apparent, especially as the number of APPN transmission groups between NNs increased (NN CPU utilization, memory requirements, and control flow bandwidth requirements are all a function of the number of NNs and transmission groups in the network).

The BX function eliminates APPN topology and APPN broadcast search flows between SNASw nodes and the SNA application hosts in the network. This feature is key to providing a reliable turnkey installation, because the network administrator no longer needs to develop in-depth knowledge of the level and characteristics of the broadcast directory search and topology update traffic in the network.


1. APPN Branch Extender Architecture Reference (SV40-0129)

SNASw was also designed to be easy to configure. The number of configuration statements and parameters are fewer than those necessary for a full-scale APPN NN deployment. Various features were added so that more often than not, many remote SNASw routers could share almost identical SNASw configuration statements.

Enterprise Extender (HPR/IP)

SNASw also supports the Enterprise Extender (EE) function. This is an optional transport facility in addition to BX support. EE offers SNA support directly over IP networks by transporting SNA traffic as connectionless User Datagram Protocol (UDP) packets (see Figure 1-2). EE support on the CS/390 host allows users to build highly reliable SNA routed networks that run natively over an IP infrastructure directly to the S/390 Enterprise Servers. EE routing at Layer 3 is performed by the IP routing infrastructure. End-to-end reliability, error recovery, flow-control, and segmentation are supported using APPN High Performance Routing (HPR) and HPR's Rapid Transport Protocol (RTP).


Figure 1-2 End-to-End EE Solution



Dependent Logical Unit Requester/Dependent Logical Unit Server

Cisco supports dependent logical unit (LU) and physical unit (PU) traffic across APPN or HPR networks through the Dependent LU Requester (DLUR) feature of SNASw (see Figure 1-3). DLUR, in conjunction with the Dependent LU Server (DLUS) function of CS/390, provides the benefits of System Services Control Point (SSCP) services and allows the session data path to be different than the one being used for the SSCP-to-PU and SSCP-to-LU control session paths. This difference allows the session to take full advantage of APPN route selection dynamics.

Figure 1-3 DLUR/DLUS Support



By putting the DLUR function in the Cisco router, remote controllers need not be upgraded to support APPN/DLUR. The dependent LUs remain visible to NetView because (CS/390) maintains awareness through DLUS-to-DLUR flows.

SNASw Benefits

Cisco SNASw contains many features and functions that allow corporations to continue their migration to an IP-based converged network that supports voice, video, and integrated data while still handling data destined for SNA applications.

Scalability

SNASw was designed specifically to support the design of highly scalable, IP-based networks that also transport SNA data traffic.

Reduced Broadcasts

SNASw, because of its inherent architecture and implementation as a BrNN node type, allows the network to scale to thousands of nodes supporting tens of thousands of RTP connections with multiple SNA sessions per RTP pipe. SNASw networks will likely contain a minimum number of NNs, and they will be located at the data center. This design eliminates significant broadcast and topology update traffic from the network.

Supports High Transaction Rates

Each SNASw router can be sized according to the transaction rate it needs to support. SNASw networks can be designed with SNASw at the remote branch, at the distribution layer, or at the data center depending on the expected transaction rates, application endpoints, and IP network design.

Transport

SNASw incorporates a number of advanced routing services and transport connectivity options.



SNASw SNA Routing Services

SNASw provides the following SNA routing functions:

- Routes SNA sessions between clients and target SNA application hosts using either APPN Intermediate Session Routing (ISR) or HPR Automatic Network Routing (ANR)
- Supports IBM standard APPN Class of Service (COS) entries and mapping of COS to IP type of service (ToS)
- Supports all types of SNA application traffic including traditional 3270 and peer LU 6.2 applications
- Supports an OS/390 Parallel Sysplex configuration, working in conjunction with the CS/390 and the IBM Workload Manager, to provide higher availability in the data center using the HPR feature
- Supports SSCP services to downstream dependent SNA devices using the DLUR feature
- Provides dynamic link connectivity using APPN connection network (CN), which eliminates much of the configuration required in networks with numerous data hosts

Responsive Mode Adaptive Rate-Based Flow Control

Early HPR implementations failed to perform well in environments subject to packet loss (for example, Frame Relay and IP transport) and performed poorly when combined with other protocols in multiprotocol networks. SNASw supports both adaptive rate-based flow control (ARB-1) and the second-generation HPR flow control architecture, called Responsive Mode ARB (ARB-2). Responsive Mode ARB addresses all the drawbacks of the earlier ARB implementation, providing faster rampup, better tolerance of lost frames, and better tolerance of multiprotocol traffic.

Native IP Data-Link Control (HPR/IP)

SNASw support for the EE function provides direct HPR over UDP connectivity. This support is used for any interface that has a configured IP address. HPR/IP can use either the real interface IP address or a loopback interface IP address as the source address for IP traffic originating from this node.

Token Ring, Ethernet, and Fiber Distributed Data Interface

SNASw natively supports connectivity to Token Ring, Ethernet, and Fiber Distributed Data Interface (FDDI) networks. In this configuration mode, the Media Access Control (MAC) address used by SNASw is the local configured or default MAC address of the interface.

Virtual Token Ring

SNASw can connect to source-route bridging (SRB) through the use of a virtual Token Ring interface. This allows the following configuration:

- Attachment to local LANs
- Connection to Frame Relay transport technologies
- Connection to CIPs and CPAs

Virtual Data Link Control

SNASw uses Virtual Data Link Control (VDLC) for two primary connectivity options:

- Transport over DLSw+ supported media
- Data link control local switching support for access to Synchronous Data Link Control (SDLC) and Qualified Logical Link Control (QLLC)

Usability, Serviceability, and Management

SNASw includes a number of features designed to improve its usability, serviceability, and management.

Dynamic Control Point Name Generation Support

When scaling the SNASw function to hundreds or thousands of nodes, many network administrators find that defining a unique control point (CP) name on each node provides unnecessary configuration overhead. Dynamic CP name generation offers the ability to use the Cisco IOS host name as the SNA CP name or to generate a CP name from an IP address. These facilities reuse one SNASw configuration across many routers and eliminate the specific configuration coordination previously required to configure a unique CP name for each SNA node in the network. However, the ability to explicitly configure each CP name still exists.

Dynamic SNA Basic Transmission Unit Size

Most SNA node implementations require specific tuning of the SNA basic transmit unit (BTU) in the configuration. SNASw analyzes the maximum transmission unit (MTU) size of the router interfaces it uses and dynamically assigns the best MTU values for that specific port. For served dependent PU 2.0 devices, SNASw uses the downstream MAXDATA value from the host and then dynamically sets the SNA BTU for that device to the MAXDATA value.

DLUR Connect-Out

SNASw can receive connect-out instructions from the IBM CS/390 (DLUs). This function allows the system to dynamically connect-out to devices that are configured on the host with the appropriate connect-out definitions. This feature allows connectivity to SNA devices in the network that were traditionally configured for connect-out from the host.

Note: DLUR connect-out can be performed over any supported data link type.

User Control of Port Limits

SNASw offers full load limiting control over the number of downstream devices (links) supported by a SNASw router. This SNASw configuration capability can control the number of devices that are serviced by this node. When the maximum number of devices is reached, SNASw no longer responds to explorers attempting to establish new connections. This allows SNASw to share the load among different SNASw nodes that offer service to the same SNA MAC addresses.

Console Message Archiving

Messages issued by SNASw are archived in a problem determination log that is queried and searched on the console or transferred to a file server for analysis. Each message has a single line that identifies the nature of the event that occurred. The log also maintains more detailed information about the message issued.

Data Link Tracing

SNA frames entering or leaving SNASw are traced to the console or to a cyclic buffer. These frames are analyzed at the router or can be transferred to a file server for later analysis. The trace can be sent to a file server in an SNA-formatted text file or in binary format readable by existing traffic analysis applications. SNASw also can capture record frames natively, eliminating the need for network analyzers to capture network events for analysis.

Interprocess Signal Tracing

The SNASw internal information is traced in binary form, offering valuable detailed internal information to Cisco support personnel. This information helps diagnose suspected defects in SNASw.



Note: This trace should only be used when requested by Cisco Technical Assistance Center (TAC) or Development Engineering support.

Trap Management Information Base Support for Advanced Network Management Awareness

SNASw supports the APPN Trap Management Information Base (MIB) (RFC 2456), which proactively sends traps with information about changes in SNA resource status. This implementation reduces the frequency of Simple Network Management Protocol (SNMP) polling that is necessary in order to manage SNA devices in the network. SNASw also supports the standard APPN MIB (RFC 2455) and the standard DLUR MIB (RFC 2232).

Positioning SNASw and DLSw+

Overview

For several years, Cisco had only one strategic technology that enabled transport of SNA over an IP backbone. That technology was DLSw+. The addition of Cisco SNASw means that Cisco provides two technologies to consolidate and transport SNA over IP.

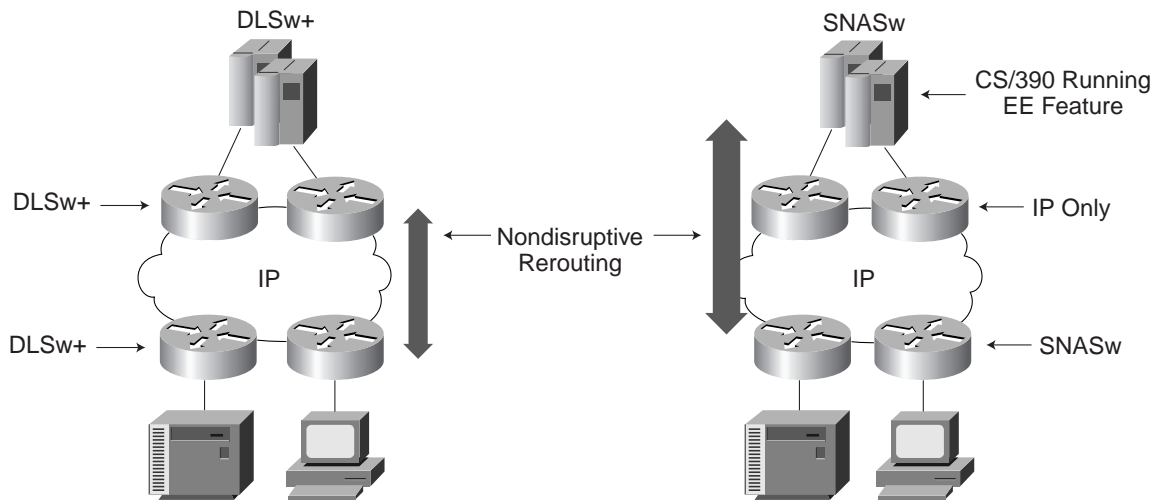
Cisco DLSw+ was created in early 1995 as a means to transport SNA over an IP network. It is based on the DLSw standards work done in the APPN Implementers Workshop (AIW) and documented in Internet Engineering Task Force (IETF) standards Request for Comment (RFC) 1795 and 2166. Cisco DLSw+ is standards compliant but includes additional features that allow it to perform better, scale larger, and provide better availability for SNA sessions.

DLSw+ enables media independence for SNA devices and simplifies SNA network configuration. It works with any SNA network: subarea, APPN, or APPN/HPR; any level of CS/390; and any LU or PU type (including PU 4). It has been a part of the Cisco IOS Software since Release 10.3 and has been continuously enhanced. To date, several hundred thousand routers run DLSw+ in production networks. Hundreds of banks, retail establishments, credit card companies, and carriers have adopted DLSw+ as their standard protocol for SNA transport over IP until now.

SNASw, Cisco's second-generation APPN platform, was created in 1999 as a replacement for Cisco's first-generation APPN NN product. It is currently available in Cisco IOS Release 12.1 and higher. APPN NN is reaching its EOE milestone concurrent with EOE for Release 11.2 on April 16, 2001, and EOE for Release 12.0 after March 2002. It is also based on standards work done in the AIW. SNASw provides two key functions: SNA routing and SNA transport in IP. SNASw uses APPN to provide SNA routing. Unlike previous APPN support in the Cisco IOS Software, SNASw includes the BX feature, an enhancement to APPN that enables it to scale. In addition, SNASw provides EE RFC 2353 support. EE enables transport of SNA data over an IP network. SNASw also provide DLUR boundary function support for dependent SNA PU 2.0 devices.

The *key* difference between the IP transport provided by SNASw EE and DLSw+ is that with SNASw EE, you can design your network to transport SNA in IP all the way into your IBM S/390 and zSeries hosts, eliminating any single points of failure in the network. SNASw EE interoperates with EE support in IBM CS/390 releases since V2R6 (with APAR OW36113) and, hence, extends the high availability characteristics of IP all the way to the EE-enabled enterprise server, as shown in Figure 2-1. With DLSw+, you can transport SNA over IP between DLSw+ peering routers, but the last hop into the mainframe is SNA.

Figure 2-1 Comparison of Availability Characteristics of DLSw+ and SNASw



The decision between these two solutions is not necessarily an either/or one. Many enterprises will deploy both. If your network currently uses FEPs for native SNA routing and you are migrating your data center from FEPs to Cisco CIP or CPA platforms (or IBM OSA-Express), then you should use SNASw somewhere in your network to provide the SNA routing function. You can deploy SNASw in the branch (as an alternative to DLSw+) or only where you currently have FEPs (in which case you can still use DLSw+ in the branches).

This chapter includes the information you need to decide which solution is appropriate for your network and where it should be deployed. It describes each technology at a high level, highlights the relevant features of each, outlines key decision criteria, describes when you need SNA routing, and suggests possible network designs.


DLSw+ Features

The DLSw+ architecture offers availability, scalability, quality of service (QoS), flexibility, and management features.

Architecture

DLSw+ transports SNA in TCP/IP (DLSw+ also offers IP-only and direct encapsulation). To monitor and maintain an end-to-end connection, DLSw+ uses a construct known as a circuit. It is comprised of two data-link control (DLC) connections and the TCP connection, as shown in Figure 2-2. TCP provides reliable delivery of data and flow control. In addition, DLSw+ has circuit-level flow control to enable traffic from one circuit to be slowed down while traffic on other circuits is unaffected.

Figure 2-2 DLSw+ Architecture



Availability

DLSw+ provides rerouting around link failures between the DLSw+ routers. If a link fails between the end system and a DLSw+ router, the sessions using that link will fail (unless the end systems are using HPR). Recovery from the failure, however, is dynamic.

Scalability

There is no known restriction on how large a hierarchical DLSw+ network you can build. The number of central site DLSw+ routers can grow in proportion to the number of remote sites and the traffic volumes—ranging from one central site router for every 100 remote locations with moderate to heavy interactive traffic to one central site router for every 200 to 300 remote branches with automatic teller or point-of-sale (POS) machine traffic. The controlling factors include the number of PUs, the number of remote sites, the traffic volumes, and the broadcast replication (typically not a problem in SNA networks).¹

QoS

DLSw+ automatically sets IP precedence, ensuring that SNA traffic receives better treatment in any network that supports IP precedence. In particular, in a Cisco environment running Cisco QoS algorithms such as Weighted Fair Queuing (WFQ) and Class-Based Weighted Fair Queuing (CBWFQ), the DLSw+ traffic is placed in a queue that is serviced more frequently. In periods of severe congestion where packet loss could potentially occur, DLSw+ packets would be among the last to be dropped. Because DLSw+ uses the same flow control method as other TCP applications, networks tuned to work well with TCP flow control will work well with DLSw+.

If DLSw+ is running in a router that is also running SNASw, DLSw+ automatically maps standard APPN COS to IP precedence ToS bits in the IP header.

Flexibility

DLSw+ is a flexible solution that addresses requirements for a variety of devices, protocols, WAN speeds, and LAN media.

1. For more information, see *DLSw+ TCP Performance* at www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/tech/dstcp_wp.htm.

Device and Protocol Support

DLSw+ supports any SNA devices or SNA networking architecture. It supports subarea, APPN, and HPR. It supports communication between any PU types including PU 4 (FEPs). Hence, DLSw+ can be used between different SNA networks transporting SNA Network Interconnection (SNI) traffic.² In addition, DLSw+ supports NetBIOS traffic.

Encapsulation Options

In addition to TCP/IP transport, DLSw+ supports UDP, Fast Sequenced Transport (FST), and direct encapsulation. FST is a high-performance encapsulation option that can be leveraged for use over higher-speed links (256 kbps or higher) when high throughput is required. FST is faster and has less overhead than TCP encapsulation, but it relies on the end systems for recovery. FST uses sequence numbers in a field in the IP header to ensure that packets arrive in order. Out-of-order packets are discarded. (DLSw+ FST transport does not support multilink transmission groups between FEPs.)

Direct encapsulation uses only link-layer encapsulation and is possible for Frame Relay or HDLC. Direct encapsulation has the least overhead but does not offer nondisruptive rerouting around link failures. In addition, it requires that the DLSw+ routers be adjacent to one another (separated only by a link or Frame Relay network), limiting network design options and forcing the DLSw+ routers to also handle WAN functionality.

Media Support

DLSw+ supports attachment to end systems over almost all media.³ However, when there are Ethernet-attached devices, network design limitations must be considered. In mixed Ethernet and Token Ring environments, the Token Ring-attached devices are limited to a frame size that will work on Ethernet. For releases prior to Cisco IOS Release 12.0(5T), in environments where Ethernet-attached devices initiate SNA connections, loops are possible. Careful design is required to prevent them. Release 12.0(5T) added a feature to enhance the redundancy characteristics when Ethernet exists in the network. See the *DLSw+ Design and Implementation Guide* at www.cisco.com/warp/public/cc/cisco/mkt/iworks/wan/dlsw/prodlit/toc_rg.htm for more information.

Management

DLSw+ can be managed with CiscoWorks Blue Maps and SNA View. SNA View provides first-level problem isolation for DLSw+, SNASw, native SNA, TN3270, and remote source-route bridging (RSRB). Maps shows a graphical view of a DLSw+ network and provides circuit-level problem determination information. **Show** commands provide additional detail, utilizing an extensive DLSw+ MIB. Hop-by-hop performance can be measured with Internetwork Performance Monitor (IPM). S/390 management of Cisco routers is available with CiscoWorks Blue Internetwork Status Monitor (ISM).

Cost

DLSw+ is a slightly less-expensive SNA transport branch solution than SNASw EE, partly because of its packaging (it is part of the IBM base feature set in Cisco IOS Software) and partly because of its size. DLSw+ branch routers require a smaller image size than SNASw routers and, hence, less memory.

2. DLSw+ transports SNI traffic between FEPs but does not replace the SNI function provided by FEPs.

3. DLSw+ has some media restrictions. For example, X.25 switched virtual circuits (SVCs) are not supported between DLSw+ and an end system. For a complete list, see the *DLSw+ Design and Implementation Guide*.

SNASw Features


The SNASw architecture also offers SNA transport over IP, availability, scalability, QoS, flexibility, and management features.

Architecture

The SNASw EE feature transports SNA using UDP/IP encapsulation. The RTP component of HPR provides reliable delivery of frames and flow control at HRP RTP EE endpoints (similar to TCP's function in a TCP/IP network). SNASw supports BX, which provides for greater scalability. The BX feature allows SNASw to provide emulated NN services to downstream SNA devices while providing an EN appearance to the upstream NN server in CS/390.

SNASw appears like an EN upstream (see Figure 2-3) and, hence, does not participate in topology updates and locates as APPN NNs do. This is what allows SNASw networks to scale. However, SNASw provides an NN image and SNA routing services to downstream SNA devices. It can register downstream devices to the CS/390 host NN central directory server and provide DLUR function so that non-APPN SNA-dependent devices can take full advantage of the transport availability afforded by HPR over IP (SNASw DLUR replaces the FEPs boundary function for SNA PU 2.0 dependent devices in subarea networks.)

Figure 2-3 SNASw Architecture



Availability

SNASw supports HPR, which provides nondisruptive rerouting around link failures between HPR RTP endpoints. Figure 2-3 shows a CS/390 enterprise server at one end and a Cisco router running SNASw at the other end. HPR could also be running in two CS/390 enterprise servers, or one CS/390 host and an EN that supports HPR.

The most common reason to implement SNASw and HPR support is to enable a router to communicate directly to IBM S/390 or zSeries hosts, encapsulating SNA in IP and enabling nondisruptive rerouting around channel or data center router outages.

Scalability

The key scalability question is how large a network can one support using SNASw. This depends upon where you implement SNASw, either in a branch servicing the resources for the branch, or in a distribution site or data center servicing the resources for multiple branches.

If SNASw is running in a branch, the limiting factor is not SNASw, but CS/390. Joint IBM and Cisco testing has demonstrated that the EE function in CS/390 scales to support at least a 2000-branch network. In these tests, SNASw EE was running in multiple routers and communicating with a single logical partition (LPAR) on the

S/390 host. Ten thousand RTP connections were established (each remote site requires multiple RTP connections, including control sessions and one RTP per COS). Thirty thousand LU-to-LU sessions across these ten thousand RTP connections were set up and maintained without any noticeable response degradation, with less than 30 percent of the S/390 CPU being utilized.

If SNASw is not running in a branch, then the scalability question is different. When SNASw runs in a distribution site or in the data center, it minimizes the number of RTP connections to CS/390. Hence CS/390 scaling is not an issue. In a distribution site or data center, SNASw also provides DLUR and NN services to the branches. The key scalability question is how many downstream resources can a single SNASw router support. To support large numbers of downstream SNA resources, multiple SNASw routers may be required. Like DLSw+, SNASw should scale to any size network. The number of central site routers will grow proportionately to the number of remote sites. The controlling factors are the number of LUs and the traffic volumes.

QoS

SNASw automatically maps standard APPN COS to IP precedence, ensuring that SNA traffic receives better treatment in any network that supports IP precedence. In particular, in a Cisco environment running Cisco IOS QoS capabilities such as CBWFQ, the high-priority SNA traffic is placed in a queue that is serviced more frequently. Low-priority SNA packets are placed in a queue that is serviced less frequently, but more frequently than a queue with unmarked packets. In periods of severe congestion, high-priority SNA packets would be among the last to be dropped.

Flexibility

SNASw has several characteristics that should be considered when assessing the appropriate implementation of the two options (DLSw+ and SNASw).

Device and Protocol Support

SNASw supports SNA only. SNASw BX support requires CS/390 to be at V2R5 or higher in your S/390 or zSeries host, and CS/390 must be running APPN. To take advantage of the SNASw EE feature (to enable SNA transport over native IP all the way into the host), you must be running CS/390 V2R6 (with APAR OW36113 applied) or higher, and CS/390 must be running APPN/HPR. Because SNASw includes DLUR support, it supports communication between PU 2.0 dependent SNA devices and the CS/390 SSCP.

Encapsulation Options

SNASw with the EE feature uses UDP encapsulated IP as the SNA transport mechanism from the SNASw EE router to the EE-enabled enterprise host server. RTP provides reliable delivery and flow control utilizing Responsive Mode adaptive rate based flow control. From the perspective of SNA APPN, EE is just another DLC connection type; to the IP network, EE is just another UDP application running in your network!

Media Support

SNASw has no media restrictions. When running SNASw in an Ethernet environment, EE provides redundancy capability and dynamic routing capabilities inherent in IP networks. Loops are not an issue as they are with DLSw+ Ethernet environments, because SNA transport between the SNASw EE router and the EE enterprise server is entirely over Layer 3 IP.



Management

SNASw can be managed with CiscoWorks Blue Maps and SNA View. SNA View provides first-level problem isolation for DLSw+, SNASw, native SNA, TN3270, and RSRB. Maps shows a graphical view of a Cisco SNASw network and provides session-level problem determination information. **Show** commands provide additional detail, utilizing an extensive SNASw MIB. Hop-by-hop performance can be measured with IPM. S/390 management of Cisco routers is available with ISM.

When SNASw EE is used end to end throughout the network, SNA traffic is essentially UDP over IP application data (the entire network is IP). Network management can therefore be performed using the Cisco IP network management tools (CiscoView 2000).

Cost

Routers running SNASw cost slightly more than routers running DLSw+ because SNASw is a separate Cisco IOS feature not included in the base IBM feature set. However, it should be noted that running SNASw EE out to remote branch locations does *not* require an SNA router in the data center, because SNASw EE transport is native IP from the remote branch SNASw router into the EE-enabled S/390 or zSeries enterprise host, thus reducing the overall cost of the SNASw EE transport solution.

Decision Criteria

For many customer environments, deciding between these two technologies is very straightforward and in some cases it is not. This section leads you through the questions you need to answer to determine if one or both of these technologies combined are appropriate in your network.

Is SNA Routing Required?

If the target SNA application is not in the same SSCP domain as the SSCP-to-PU and SSCP-to-LU control sessions, SNA routing is required. Without SNASw or APPN, this routing is done by the owning SSCP itself, forcing the data center hosts to perform SNA routing and raising the application cost by increasing mainframe CPU utilization. Therefore, if SNA routing is required, you need SNASw. SNA routing, which is required in any cross-domain environment, has traditionally been done by FEPs, but the majority of customer enterprises today have chosen to migrate from FEPs to higher-speed, multiprotocol routers to save money and to position their networks for multiprotocol data/voice/video applications supported by the Cisco Architecture for Voice, Video, and Integrated Data (AVVID), as well as applications such as IBM WebSphere on S/390 and zSeries Parallel Sysplex mainframes. If you want to replace your FEPs with routers, and your FEPs are currently routing SNA traffic between different LPARs, then you will most likely want your Cisco routers to provide that same functionality. SNASw provides that capability. You may still need to decide where you want to run SNASw—that is, at the branch or at a distribution site. That decision is covered in the “Network Design” section.

Are You Ready?

If you want to consolidate SNA and IP traffic onto a single backbone, then you can use either DLSw+, SNASw, or DLSw+ and SNASw combined. Before you weigh the technical merits of each, answer these simple questions:

- *What operating system level are you running on your S/390 or mainframe?*—The SNASw BX feature requires CS/390 V2R5 or higher. The EE feature in CS/390 requires OS/390 V2R6 (with APAR OW44611 applied) or higher.
- *What Cisco IOS release are you running on your branch routers?*—SNASw requires Cisco IOS Release 12.1 or higher.
- *Are you running APPN or APPN/HPR in CS/390 today?*—The biggest advantages SNASw brings is the ability to provide necessary SNA routing (BX feature) and transport SNA in native IP all the way into your S/390 (EE feature). However, to take advantage of those capabilities, you need to be CS/390 APPN-enabled (for BX) and APPN/HPR-enabled (for EE) and be at the required version and release of CS/390 software.

Additional considerations include availability, cost, and application direction. SNASw offers the highest potential availability (although it requires much newer S/390 and router software). It costs slightly more per branch router than DLSw+ because of the additional cost of the SNASw Cisco IOS feature set. If you intend to continue supporting existing mission-critical SNA applications on your S/390 and zSeries mainframes for several more years while developing newer applications in IP, you should consider SNASw because it can bring you the highest availability possible. You may want to start with SNASw only in the data center (combined with DLSw+ in data center peering routers), and then migrate SNASw EE out to distribution sites or remote branch offices. See the “Network Design” section for more detail.

Network Design

The technologies described in this chapter are not mutually exclusive. Both SNASw and DLSw+ can be used in the same network, and in fact, a large number of Cisco customer implementations of have done just that.

In many customer situations, DLSw+ supports SNA WAN transport to remote DLSw+ peering routers at remote branch locations. SNASw can be deployed in locations where the network has FEPs to replace necessary SNA routing previously provided by the FEPs, and to replace FEP boundary function support for dependent SNA devices. If FEPs are in the data center or distribution sites, that is where the SNASw routers reside. At the data center or distribution site, SNASw can run in the same routers that currently handle DLSw+. This implementation allows organizations to isolate the SNASw function to an “SNA router”—so named because it provides the functions required to support necessary SNA routing of client sessions directly to the target application host in addition to providing DLSw+ peer termination points for WAN transport of SNA from remote SNA devices.


However, there are many network design possibilities. Rather than go through them all, this chapter discusses three hypothetical case studies, explains the chosen design, and describes the reason for the choices. Because the focus of this chapter is to describe DLSw+ and SNASw as SNA-over-IP transport alternatives, none of these case studies shows SNASw being deployed without the SNASw EE feature (although it is certainly possible to do so predicated on the environment and requirements).

Case Study 1—No HPR over IP Transport (DLSw+ Only)

In this case study, the enterprise has no requirement for APPN because it has no cross-domain sessions, is not using FEPs to route cross-domain traffic, or routes cross-domain traffic through CS/390 enterprise servers. The enterprise is not planning to change its SNA applications and is planning to maintain the status quo in the data center for some time under the reasoning “it works; therefore, don’t touch!”

The simplest, surest, and lowest-cost solution for this customer is to use DLSw+ to transport SNA traffic over an IP backbone back to a data center router. The customer chose to keep the SNA data center router separate from the WAN distribution router to simplify change management and maximize availability. The SNA data center router runs the Cisco IOS level that has all the DLSw+ features needed, and the customer does not want to modify it to pick up the latest compression or security feature (and vice versa for the WAN distribution router). Two CIPs (one primary and one backup) run Cisco SNA (CSNA) to handle all the SNA traffic, and four Cisco 7200 Series routers run DLSw+ (to handle a 600-branch network), including one DLSw+ router used only for backup. Figure 2-4 shows this design.

Figure 2-4 DLSw+ Design



Case Study 2—DLSw+ and SNASw


In this case study, the enterprise wants to leverage its Parallel Sysplex complex and achieve the high availability it affords. The customer is migrating to Cisco Catalyst® 6500 Gigabit Ethernet switches with Multilayer Switch Feature Cards (MSFC) installed for IP Layer 3 support in the data center, and new host applications are being written to run IP natively. However, it will be several years before a complete migration from SNA to IP applications is complete, and in the interim, the customer wants the high availability and design simplicity afforded by having an all-IP data center today.

This enterprise already uses DLSw+ to transport SNA traffic over an IP backbone. The customer chose not to deploy SNASw EE out to the remote branches at this point because the DLSw+ network has been in place and quite stable for a long time (if network outages occur over the DLSw+ network, they affect only a small portion of the network and are recovered automatically). However, the customer wants to ensure that a CIP or channel outage (which today would bring down almost the entire network) can be handled transparently and nondisruptively. Hence, the customer is adding SNASw to hub-end data center routers terminating DLSw+ peer connections coming in from remote branch DLSw+ routers. The BX capability of SNASw is providing necessary SNA routing for downstream SNA devices, while at the same time the SNASw EE feature transports SNA traffic natively over IP into the CS/390 enterprise server upstream. By doing this, the customer has eliminated Token Ring source-route bridge requirements for maintaining multiple active redundant bridged routing information field (RIF) paths to the mainframe (required for Logical Link Control, type 2 [LLC2] bridged transport), and has also eliminated the potential for loop problems that can occur in a bridged Ethernet environment. In addition, should a channel failure occur, IP immediately reroutes traffic and SNA sessions are not impacted (since EE uses HPR for nondisruptive session path switching). Finally, this design positions the customer to use Cisco Catalyst 6500 Gigabit Ethernet switches connected to IBM S/390 OSA-Express for SNA traffic transport over IP (HPR/IP).⁴

4. The OSA-Express Gigabit Ethernet card is for TCP/IP environments only. This card only supports SNA traffic when SNA is encapsulated in IP using the EE support in OS/390 Version 2 Release 6 (with APAR OW44611) or higher.

As in the previous case study, this enterprise chose to keep the SNA data center router separate from the WAN distribution router to simplify change management and maximize availability. Two CIPs (one primary and one backup) run IP to handle all the SNA traffic, and six Cisco 7200 Series routers run DLSw+ and SNASw (to handle a 1000-branch network), including one DLSw+/SNASw router used only for backup. Figure 2-5 shows this design.

Figure 2-5 Combined SNASw and DLSw+ Design




Case Study 3—SNASw Only

In this case study, the enterprise demands the highest availability for its SNA applications. The customer has invested a great deal in rewriting applications to LU 6.2 and wants to continue to leverage that SNA application investment. The customer was running separate networks for SNA and IP and decided to consolidate using SNASw EE HPR over IP support for SNA transport natively over IP. The network is already at the latest operating system level and is running APPN/HPR and EE support in CS/390.

The customer has 200 regional offices that run SNASw EE. From the branch into the S/390 host, the SNA traffic is transported in IP. Hence, there is no need for SNA routers in the data center. The customer leverages the Cisco QoS features to ensure that the interactive SNA and Telnet traffic take precedence over SNA batch and FTP traffic. Figure 2-6 shows this design.

Figure 2-6 SNASw Design



Conclusion

Cisco DLSw+ is a proven method of transporting SNA traffic over an IP network. The downside of DLSw+ is that the final hop into the mainframe is SNA. SNASw EE can eliminate this issue by combining the best qualities of APPN and DLSw+. SNASw is going to play a key role in many data centers in the future. Whether it is right for you today depends on your application direction, cost constraints, OS/390 level, and CS/390 configuration. If you are moving to APPN and APPN/HPR, SNASw can provide enhanced availability and data center flexibility. DLSw+ will also continue to play a key role as a WAN transport for SNA traffic for the foreseeable future.

Implementation of SNASw

Getting Started

This chapter describes some of the general considerations for SNASw BX, EE, and DLUR deployment; minimum required SNASw configuration and considerations; SNASw advanced configuration features; and the primary scenarios for implementing SNASw. This section assumes that you are somewhat familiar with basic Cisco router configuration. The chapter will extensively discuss general deployment and network design considerations for optimally deploying SNASw in your network.

When Is APPN Required?

The first question you need to answer is what do you need. There are cases where APPN is required and cases where it is not. If you have multiple mainframes or LPARs in your data center, you will need to make an SNA application routing decision somewhere in your network. If you do not perform that SNA routing decision in the network, then it will get done in the host enterprise servers, causing SNA LU-to-LU application session routing to traverse through the S/390 communication management configuration (CMC) hosts owning the SSCP-to-PU and SSCP-to-LU control sessions.

There are only two possible SNA routers that can route SNA client sessions directly to target application hosts: IBM FEPs and Cisco SNASw routers. Traditionally, SNA routing has been done on FEPs, but the majority of companies today have chosen to migrate from FEPs to high-speed, multiprotocol router-based networks to save money and to position their enterprises for converged IP infrastructure applications such as Cisco AVVID and IBM Parallel Sysplex and WebSphere on IBM S/390 and zSeries enterprise servers.

If SNA routing is required and FEPs are installed, you must make a decision either to keep the FEPs or to replace them with Cisco SNASw routers. If you are replacing FEPs with Cisco CIP or CPA channel-attached routers or Catalyst 6500 Gigabit Ethernet switches attached to S/390 (or zSeries) hosts and plan to continue supporting SNA application routing in your environment, you will need to implement SNASw somewhere in your network.

Although SNASw can be extremely beneficial in your consolidated network IP infrastructure, it is important to carefully plan and design how and where SNASw is deployed. It is not necessary to provide SNA routing throughout the entire network. In fact, a single SNA routing decision is all that is actually required.

In many situations, the SNA routing decision and boundary function support can occur in the data center, eliminating native SNA routing from existing APPN NN-enabled devices cascaded downstream in the network. This design is especially effective in situations where existing large-scale DLSw+ and APPN NN networks are in place. However, moving that SNA routing decision and boundary function support to the aggregation layer or

regional office locations may be highly beneficial if traffic is consistently routed between multiple data centers (when using non-DLSw+ SNA encapsulation for IP, the SNA boundary should be placed at edges of the network whenever possible).

APPN, HPR, DLUR, EE (HPR/IP), and BX are important in developing the IT infrastructure for e-business and therefore are significant functions of SNASw. Most mission-critical business information comes from an SNA heritage and, even today, a large percentage of such traffic is based on SNA applications. The IBM S/390 and zSeries Parallel Sysplex mainframes have effectively incorporated dynamics, QoS, scalability, and workload-balancing functions using APPN and HPR. To efficiently access SNA data in this environment, you should use HPR inside the Parallel Sysplex and Cisco SNASw somewhere in the network for making SNA routing decisions for peripheral SNA devices.

SNASw General Considerations

There are several considerations that you need to take into account when considering a migration to SNASw:

- SNASw does not support APPN NNs or APPN peripheral border nodes (PBN) connected downstream from it. Only downstream APPN ENs, LENs, and dependent SNA PU 2.0 devices (using SNASw DLUR support) connected downstream from SNASw are supported.
- SNASw does not support CS/390 (VTAM) hosts of any APPN node type connected downstream from an SNASw router.
- SNASw is not an SNI replacement solution. APPN extended border node (EBN) support for HPR over IP (EE) connections between CS/390 hosts is a potential replacement for SNI (SNASw has no role in the EBN environment).
- SNASw does not support DLUR routers cascaded downstream from another DLUR. The IBM DLUR/DLUS architecture (supported by SNASw) has a restriction that if a DLUR is required, it can only be implemented in the DLUR *directly connected* to the upstream DLUS NN server host.
- SNASw BX routers can be cascaded downstream from other SNASw BX routers for support of independent LUs (LU 6.2) traffic, but this is not a recommended best-practice network design deployment for SNASw. The objective for designing scalable enterprise SNA over IP networks is to eliminate all instances of intermediate SNA session routing by deploying a single-level SNA routing network model.

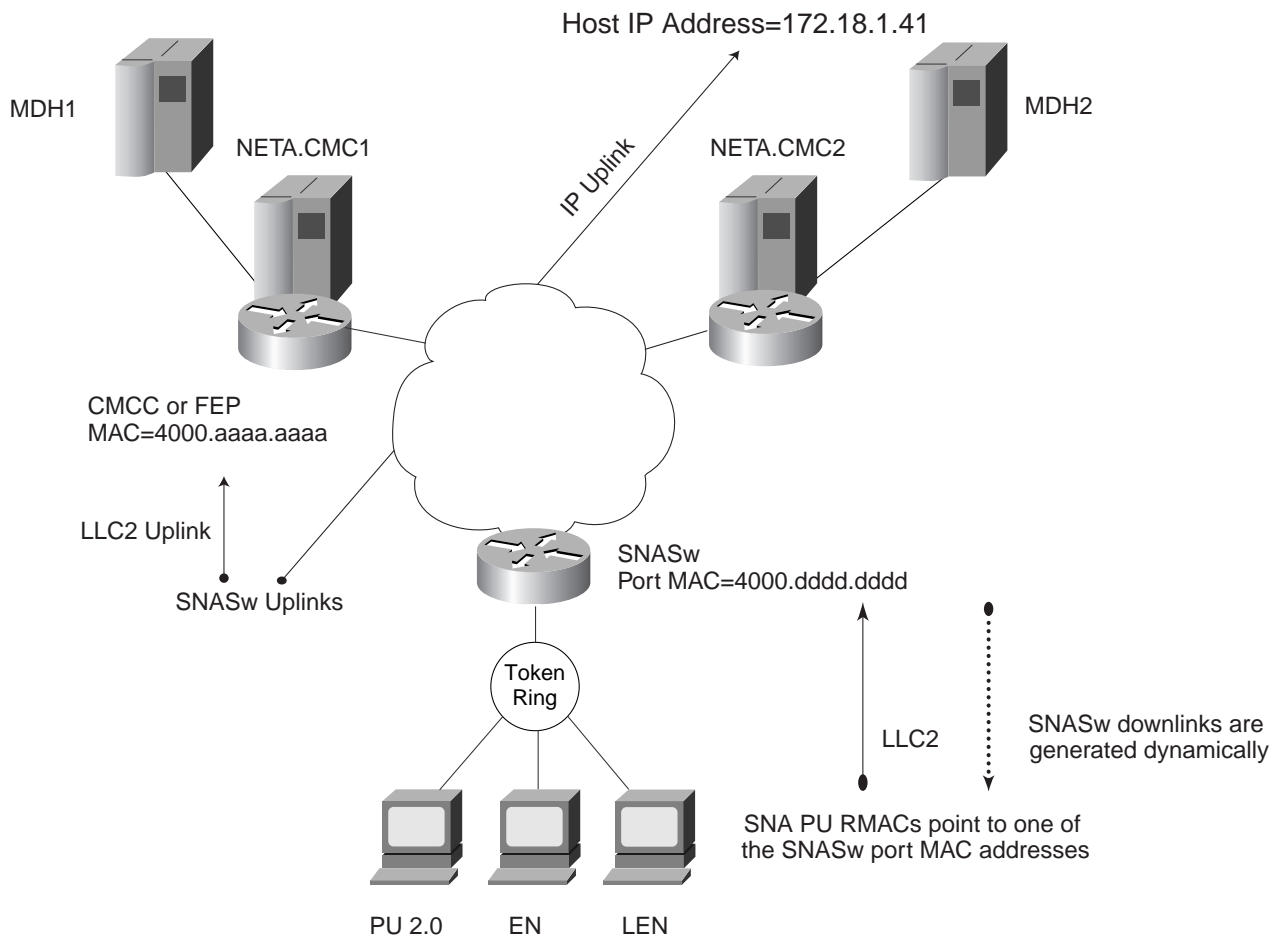
Required SNASw Configuration

To enable SNASw in a Cisco router, you must configure the following in this order as illustrated in Figure 3-1:

- SNASw CP
- SNASw port
- SNASw upstream link to the NN server/DLUS host (and backup NN server/DLUS)


Note: SNASw supersedes all functionality previously available in the APPN NN feature in the Cisco IOS Software. SNASw configuration does not accept the previous APPN configuration commands.

Figure 3-1 Required Configuration



This order is specified because of the hierarchical nature of SNASw definitions. If your network consists of many APPN application EN hosts that communicate with each other, then configure the SNASw CP, SNASw port, and SNASw link to upstream NN servers (and backup NN servers) and use SNASw connection network support to dynamically activate the other EN-to-EN links, as shown in Figure 3-2. (SNASw connection network is discussed extensively later in this chapter.)

Figure 3-2 SNASw and Connection Network



Every APPN node requires an SNASw CP definition, which uniquely identifies the node within a network. Configuring a CP name dynamically activates SNASw. Removing a CP name definition deactivates it.

Note: You should configure a unique CP name for SNASw instead of using the same one you used for APPN NN (PSNA). Not using a unique name for SNASw could result in problems due to the fact that the existing NN topology may not have purged the APPN NN CP name from the topology databases.


Every interface that is used for SNASw communication also requires an SNASw port definition statement. An SNASw port definition associates SNA capabilities with a specific interface that SNASw uses for transport. The port associates with a Token Ring, Ethernet, FDDI, or ATM interface that points downstream. SNASw also enables support for DLSw+, SDLC, QLLC, and SRB transport downstream. The port also defines an interface with a MAC address, which serves as the destination MAC address (DMAC) to which downstream SNA devices connect. SNASw ports dynamically start when they are configured in the router unless the **nostart** keyword is configured on the SNASw port definition.

You must define the interface over which the port communicates before you define and activate the SNASw port. SNASw ports do not dynamically adjust to interface configuration changes that are made when SNASw is active. For example, if you change an interface MAC address or MTU size, SNASw may not recognize the new value. If you want to make changes to an interface and want SNASw to adjust to the changes, you may need to either delete and redefine the port that is using that interface or stop and restart SNASw. To manually activate an SNASw port, issue the **snasw start port portname** command. To manually deactivate an SNASw port, issue the **snasw stop port portname** command. Details of the various transport options supported by SNASw are discussed later in this chapter.

SNASw communicates with upstream devices over a link object (uplink). In all SNASw designs, you must always explicitly predefine the SNASw host uplink for the CP-to-CP session between the SNASw CP and upstream NN server CP, which typically is also serving as the upstream DLUS (see Figure 3-3). Without at least

one NN server uplink, SNASw is unable to provide connectivity to other application EN (or LEN) devices, or to provide SSCP services to downstream dependent devices supported by the SNASw DLUR capability. Therefore, at least one uplink definition is typical in every SNASw network. It is also a common design practice to have an additional uplink to another upstream NN server/DLUS such that if the primary NN server is down, the CP-to-CP session between SNASw and the server remains fully functional on the backup.

Figure 3-3 SNASw Uplinks and Downlinks



In addition, you can also define SNASw link definitions to other target EN application hosts if desired, but using SNASw connection network support to minimize the number of links (uplinks) is by far the more preferred method (SNASw connection network support is discussed in a later section of this chapter). SNASw can support a maximum of 10-12 predefined host uplinks configured in the router. An upstream link is not required if a partner node initiates the connection, because the link definition is built dynamically when the partner node initiates it.

For all links requiring configuration in the SNASw router (such as the links to upstream NN server and to interchange nodes [ICNs] in situations as described in the section SNASw Connection Network Support later in this chapter), configure them to point to either a remote MAC address such as a CIP or CPA MAC address (for LLC transport) or an IP address on the host (for HPR/IP EE transport). This identifies the partner address to which SNASw attempts to establish the link. SNASw ports dynamically start when they are configured unless the `nostart` keyword is configured on the `snasw` port definition.

Do not use the `snasw link` command to connect to client workstations and devices downstream being serviced by the SNASw router (as was illustrated previously in Figure 3-1). Downstream SNA devices should be configured with an outbound connection to the MAC address of the active SNASw port servicing downstream devices on the SNASw router. However, there are two potentially useful options you can configure on SNASw for downstream devices:

- You can limit the maximum number of link station connections into an SNASw router from downstream devices attempting inbound connections to SNASw. Enable this function by configuring the **max-links link limit value** option in the **snasw port** command. This option provides the ability to load limit the number of downstream SNA device connections into an SNASw BX router. Multiple SNASw routers with duplicate MAC address endpoints servicing downstream devices can then be utilized to load limit connections into SNASw routers across multiple SNASw router MAC address endpoints.
- You can define the location of a specific resource (which is required for LEN type devices) by configuring the **snasw location** command. Use this function when a LEN link is established with a partner node. The command allows SNASw to route session requests over the LEN link to the resources named in the **snasw location** statement.

The configuration of **snasw location** is not required in all LEN resource situations (you never need to define **snasw location** statements for dependent LUs). For example, in the case of independent LUs if the LEN node device always initiates the first session, or if the LEN CP name matches the names used for the independent LU-to-LU sessions, **snasw location** definitions are not required.

Note: For more details regarding SNASw commands, see the “SNA Switching Services Commands” chapter in *Cisco IOS Bridging and IBM Networking Command Reference, Volume II*. For more information about SNASw configuration guidelines, see the “SNA Switching Services Configuration Guide” chapter in *Cisco IOS Bridging and IBM Networking Configuration Guide*.

Supported SNASw Link Types and Transport Options

The SNASw subsystem supports a wide range of link types to establish SNA connections with upstream and downstream devices. Supported link types and interfaces include the following:

- Native SNA transport on Token Ring, Ethernet, and FDDI
- Virtual Token Ring interfaces that support source-route bridged connections to local LANs and channel interface cards such as the CIP and CPA
- SNA over Frame Relay using bridged Layer 2 format RFC 1490 frames (Boundary Network Node [BNN] and Boundary Access Node [BAN])
- DLSw+ transport using VDLC
- Attachment to SDLC and QLLC links using DLSw+ local switching support

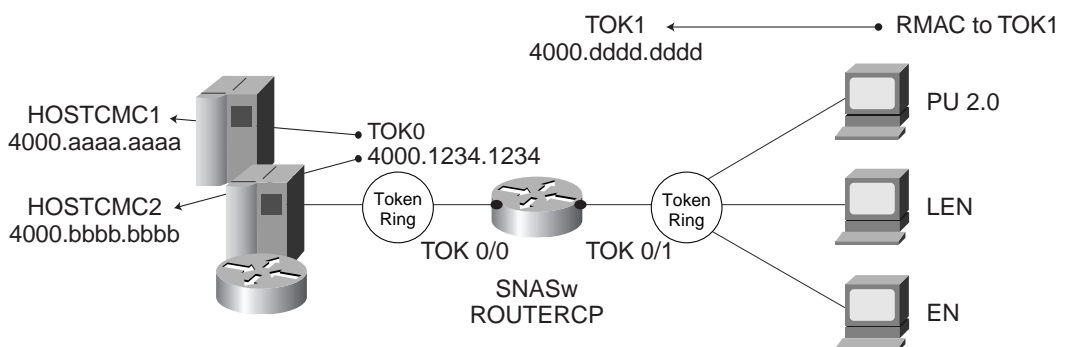
SNASw does not support the following transport options:

- SNA over Frame Relay using routed format RFC 1490 frames (BNN)
- ATM RFC 1483
- Point-to-Point Protocol (PPP) support
- Direct connections to SDLC and QLLC (SDLC and QLLC connections into SNASw are supported by DLSw+ local switching support using VDLC)

Native LAN Transport

SNASw natively supports connectivity to Token Ring (as illustrated in Figure 3-4), Ethernet, and FDDI networks. In this configuration mode, the MAC address used by the SNASw port is the locally configured or default MAC address of the physical interface on the SNASw router.

Figure 3-4 Native LAN Transport



```
interface TokenRing 0/0
mac-address 4000.1234.1234
ring-speed 16
interface TokenRing0/1
mac-address 4000.ddd.ddd
ring-speed 16
```

SNASw/APPN Control Point Name
snasw cpname NETA.ROUTERCP

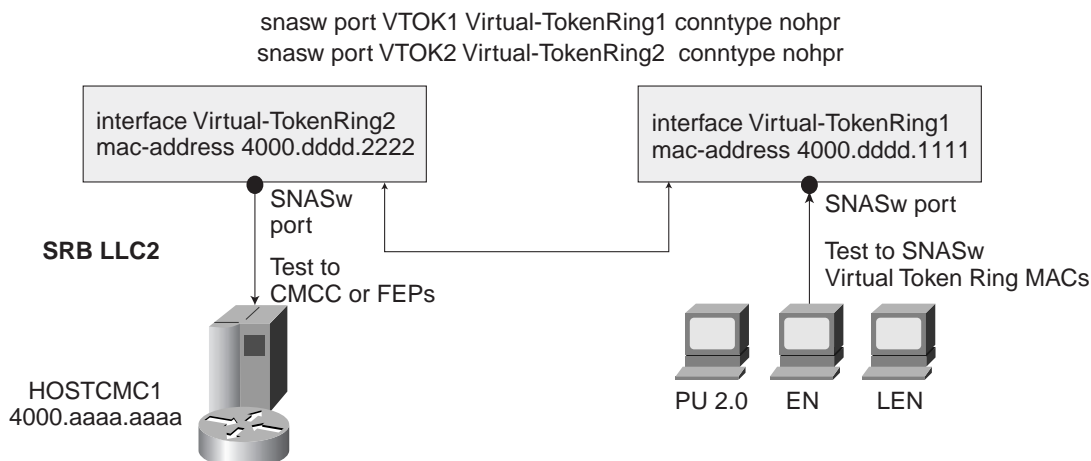
SNASw Port for Downstream SNA Devices
snasw port TOK1 TokenRing0/1 conntype nohpr

SNASw Port and Links for Upstream Hosts
snasw port TOK0 TokenRing0/0 conntype nohpr
snasw link HOSTCMC1 port TOK0 rmac 4000.aaaa.aaaa
snasw link HOSTCMC2 port TOK0 rmac 4000.bbbb.bbbb

Virtual Token Ring Transport

Virtual Token Ring and SRB allows SNASw to respond to multiple MAC address endpoints mapped to a single physical LAN interface on a SNASw router (as shown in Figure 3-5). Because there is no limit to the number of virtual Token Ring interfaces you can configure in the router, multiple virtual Token Ring interface MAC addresses can respond to downstream device SNA requests over the same LAN interface (when using native LAN support, SNASw responds to requests to the target MAC address configured on the local LAN interface only). This can be very beneficial when migrating from multiple IBM FEPs to Cisco CIPs or CPAs and deploying SNASw to replace SNA routing functionality. Each FEP Token Ring interface coupler (TIC) MAC address previously configured on the FEP can be replicated on individual virtual Token Ring interfaces configured on the SNASw router. Virtual Token Ring and SRB can also be used to connect (bridge) SNASw traffic to upstream hosts using LLC transport over the CIP or CPA.

Figure 3-5 Virtual Token Ring Transport



Virtual Token Ring and SRB can also connect SNASw to an SNA Frame Relay Layer 2 bridged infrastructure. Frame Relay Access Support (FRAS) host and SRB Frame Relay are configured to connect virtual Token Ring interfaces that offer SNASw support for Frame Relay BAN or BNN technology.

VDLC Transport

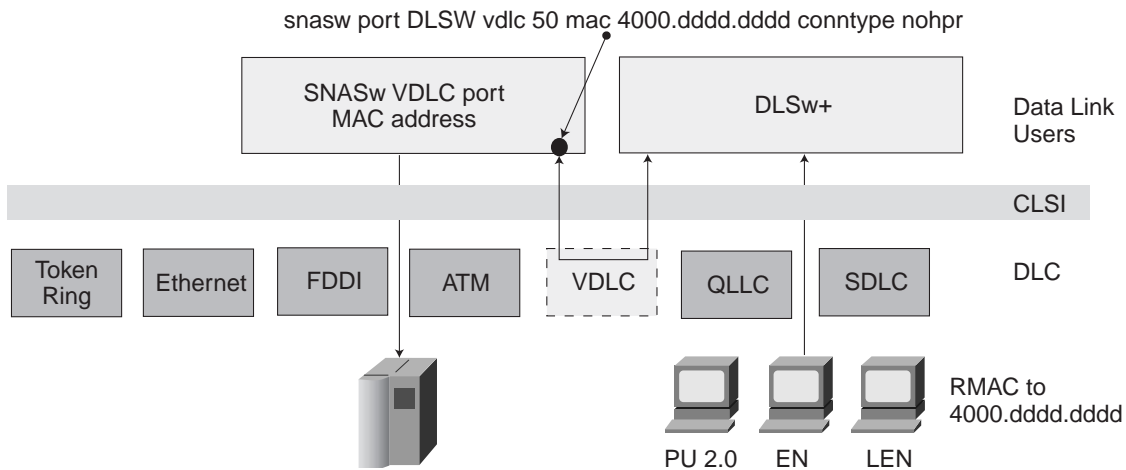
SNASw uses VDLC to connect to DLSw+ transport and local switching technologies. VDLC is used for a number of connectivity options, including the following:

- Transport over DLSw+ supported media
- DLC local switching support for access to SDLC and QLLC

Using VDLC, SNASw gains full access to DLSw+ SNA transport capabilities, including DLSw+ transport over IP networks, DLSw+ transport over direct interfaces, and DLSw+ support of direct Frame Relay encapsulation (without using IP). SNASw also gains access to devices connecting through SDLC and QLLC (see Figure 3-6).

Note: SDLC and QLLC are transparent to the SNASw code.

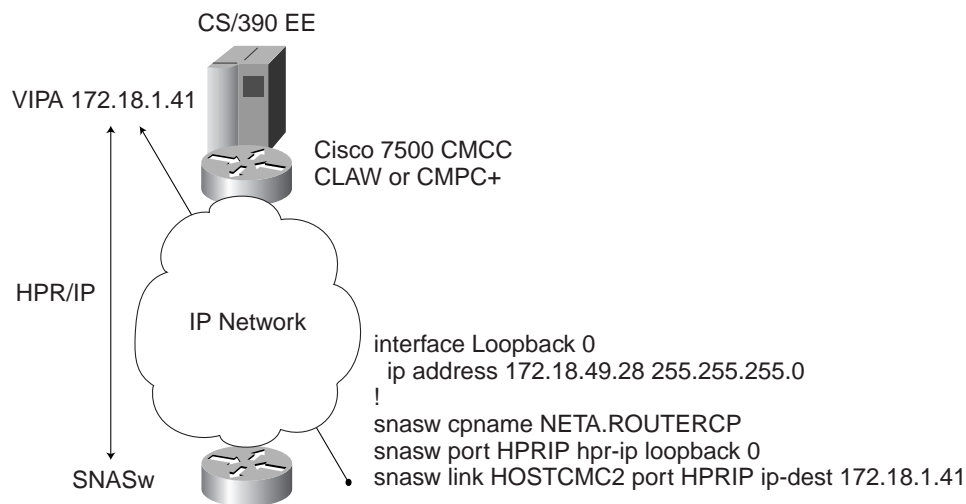
Figure 3-6 VDLC Transport



Native IP DLC Transport

SNASw support for the EE function provides direct HPR over IP/UDP connectivity for SNA host transport. This support is configured for any interface that has a configured IP address. HPR/IP uses the interface IP address (such as the loopback interface) as the source IP address for IP traffic originating from this node (see Figure 3-7).

Figure 3-7 Native IP DLC Transport



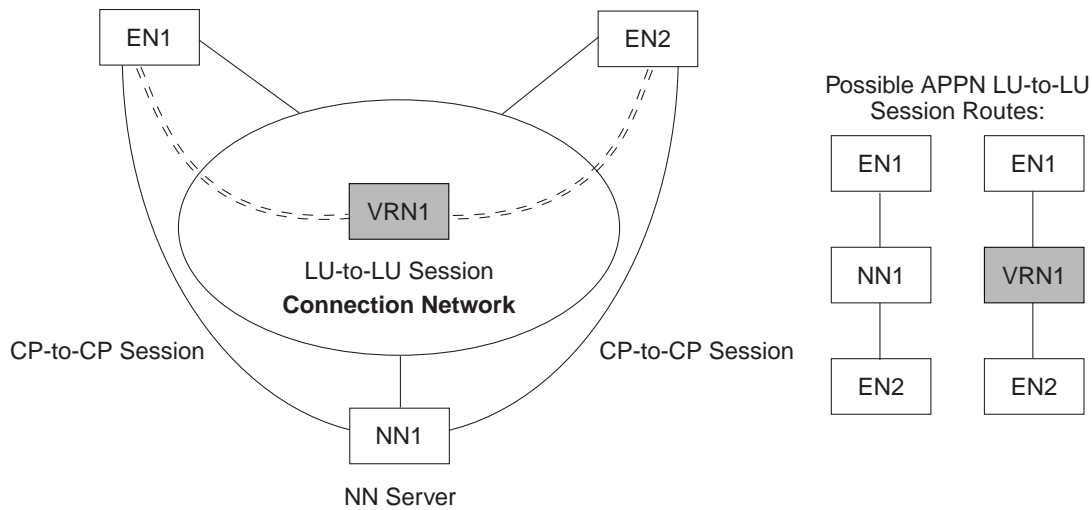
SNASw Connection Network Support

The true value of SNASw is that it connects SNA clients originating SNA sessions to the target application data host directly without having the session path traverse through hosts or CMCs that originated the SSCP-to-PU and SSCP-to-LU control sessions.

The method traditionally used by many APPN installations for enabling host connections is the configuration of individual APPN links (uplinks) to each application EN host. In networks with just a few upstream hosts and SNASw routers installed, explicitly defining the upstream host links is typically not a problem (SNASw supports approximately 10-12 host uplinks). However, most enterprise SNA networks have a much larger number of application EN hosts, making the task of defining host links more cumbersome and defeating the purpose. Because most actual user sessions terminate in applications running on application hosts in data center, distribution layer, or remote branch locations, the best practice network implementation approach is to have direct *dynamic* links between SNASw routers and the application EN hosts themselves.

The application of SNASw connection network allows a simple definition of a common virtual routing node (VRN) to which all ENs and SNASw routers can connect, as illustrated in Figure 3-8 (connection network support is configured on the SNASw port definition). In a connection network environment, hosts and SNASw nodes are configured to belong to the same VRN. The actual terminology that is often used to refer to this transport infrastructure is shared access transport facility (SATF). When an APPN EN (including an SNASw router) registers with its NN server, its VRN (if defined in the SNASw port configuration) is recorded. When that node subsequently requests a session, the NN server compares the VRNs of the requesting node and the destination node and, if they are a match, provides to the requester the direct route path to the destination. This allows LU-to-LU sessions to be established directly and dynamically between ENs, reducing network latency to a minimum and freeing NN resources for other work.

Figure 3-8 VRNs and Connection Networks



With base APPN ISR routing, connection network can be implemented over LLC2 connectivity between APPN ENs (at Layer 2). Therefore, the potential for connectivity exists when a LAN infrastructure is available to the APPN nodes requiring links and there is a means of bridging these LANs together. In a data center environment this is usually a source-route spanning tree implementation, while for geographically dispersed remote sites Cisco DLSw+ provides a robust and manageable means of transporting the underlying LLC2 traffic over a TCP/IP network (using DLSw+ for SNA transport) to SNASw/DLSw+ central site or regional aggregation routers.

With SNASw EE and HPR/IP ANR routing, however, the IP network *itself* becomes the connection network at Layer 3. The common VRN represents the existence of a link into the common IP network. When SNA sessions need to be established between an SNASw router and target application host, APPN's topology and routing services component recognizes the existence of this common VRN link and causes a direct EE link to be dynamically set up between the two ENs over the IP network.

You need to be aware of a couple issues with connection network if you plan to have EE (HPR-only) connections adjacent to interchange transmission groups, and if any of the hosts connected are ICNs. A common scenario for this situation is where an OS/390 and CS/390 APPN EE host is also an SNI gateway to another network (ICN). Before OS/390 and CS/390 V2R10 (and releases before IBM APAR OW44611), this did not allow sessions to cross-domain subarea partners to exit an ICN via an HPR connection unless the connection was only one hop away from the target EN.


Even with OS/390 and CS/390 V2R10 (or higher) and the fix to IBM APAR OW44611 applied, there will still be one scenario that will not support HPR on an APPN link immediately adjacent to an interchange transmission group. This is when an ICN defines a connection to a connection network (VRN) and a session is attempted from the subarea network through the ICN and then into APPN over the VRN. Refer to IBM APAR OW44611 for more information regarding this limitation.

SNASw Connection Network Support for Branch-to-Branch Traffic

SNASw connection network also effectively addresses the issues of building partial or fully meshed SNASw networks when EN resources downstream from SNASw BX routers in different remote locations need to communicate and establish LU-to-LU sessions. For example, EN1 in Figure 3-9 under SNASw Router A needs

to communicate with EN2 under SNASw Router B. By defining SNASw Router A and SNASw Router B to be part of the same VRN (VRN A in this example), direct LU-to-LU session traffic between EN1 and EN2 is supported.

Figure 3-9 Connection Network Branch-to-Branch Example




IBM CS/390 Connection Network Support

For information regarding the implementation of connection network support in IBM CS/390, refer to the *OS/390 IBM CS SNA Network Implementation Guide* (IBM Publication SC31-8563) and to Appendix C, Enterprise Extender (HPR/IP) Sample Configuration.

Enabling SNASw DLUR Support

DLUR and DLUS are APPN architecture features that allow dependent LU (DLU) PU 2.0 traffic to flow over an APPN network. Before the introduction of the DLUR/DLUS feature, the APPN architecture assumed that all nodes in a network could initiate peer-to-peer traffic. However, DLUs cannot do this, because it requires the SSCP to notify the application and then send the BIND request. Initiating the legacy sessions requires a client/server relationship between the Cisco SNASw router (which natively supports DLUR) and the host SSCP enabled as a DLUS. A pair of LU 6.2 session pipes is established between the SNASw DLUR router and the DLUS (one session is established by each endpoint in both directions). These sessions are then used to transport the SSCP-to-PU and SSCP-to-LU control messages that must flow in order to activate the legacy resources and to initiate their LU-to-LU sessions directly to the target application host (see Figure 3-10).

Figure 3-10 SNASw DLUR Support




The following steps illustrate how a DLUR/DLUS session is performed:

- Step 1. The host must send an activate LU (ACTLU) message to the LU to activate a DLU. Because this message is not recognized and supported natively in an APPN environment, it is carried as encapsulated data on the LU 6.2 session pipes.
- Step 2. DLUR then unencapsulates it and passes it to the legacy LU.
- Step 3. The DLU session request is passed to the CS/390 NN/DLUS, where it is processed as legacy traffic.
- Step 4. DLUS sends a message to the application host, which is responsible for sending the BIND.
- Step 5. SNASw establishes the LU-to-LU session directly to the target application host.

When SNASw is configured and enabled in a Cisco router, DLUR functionality is automatically enabled by default and does not require any configuration effort whatsoever. This is advantageous because it allows the SNASw DLUR to connect dynamically to an upstream DLUS over the active CP-to-CP session between the SNASw router and the upstream NN server host. If the CP-to-CP session fails and SNASw re-establishes the CP-to-CP session with another (backup) upstream NN server, SNASw DLUR automatically reconnects to the DLUS on the backup NN server if DLUS functionality is enabled on the backup host.

For example, if you have five upstream host links defined to NN servers from SNASw, SNASw uses the NN to which it has established the CP-to-CP session as its DLUS. The other four upstream NN links can provide backup DLUS services if the primary DLUS fails. This scenario provides for five possible DLUS backups servers for the SNASw DLUR (as opposed to only one primary and one backup DLUS server when you hardcode the DLUR in SNASw as is described in the next paragraph).

SNASw does provide the ability to explicitly configure the primary as well as backup DLUS server upstream using the `snasw dlus` configuration command (DLUR is enabled by default in a Cisco router). If you wish to manually configure DLUR/DLUS support with SNASw, you must specify the fully qualified name of the primary



DLUS (**snasw dlus primary-dlus-name**). The following additional DLUR configuration options are available for DLUR/DLUS support with SNASw:

- Specify a backup DLUS by configuring the **backup backup dlus name** option. Define a backup DLUS that activates when the primary DLUS is unreachable or unable to service a downstream device.
- Define exclusivity for inbound connections to primary DLUS by configuring the **prefer-active** option.
 - If an active DLUR/DLUS connection was established, you can specify exclusivity on the active DLUS connection so that an incoming PU cannot connect to another DLUS.
 - If you do not specify the **prefer-active** keyword, each downstream connected station attempts connections to both the primary and backup DLUS until the device receives DLUS services.

SNASw defaults to using its current active upstream NN server as the preferred DLUS for the node. To override this default and explicitly configure the DLUS name, configure the **snasw dlus** command.

In addition, you can configure node-wide defaults for the DLUS and backup DLUS that the SNASw DLUR router contacts:

- Define the number of connection attempts by configuring the **retry interval** option.
 - You can specify the interval between connection attempts to a DLUS (except when serving an exclusive PU).
 - If you specify an interval, then you must specify the **retry count** option. This option specifies the number of connection attempts the DLUR makes to the current or primary DLUS before connecting to a backup or currently nonactive DLUS.

Customer Scenarios for SNASw Deployment

This section presents three scenarios for customer migration and deployment of SNASw:

- Customers migrating from APPN NN to SNASw
- Customers migrating from FEPs to SNASw
- Customers migrating from Token Ring to Ethernet

Customers Migrating from APPN NN to SNASw

There are a number of problems associated with building large APPN ISR and HPR networks composed of large numbers of APPN NNs. These issues are described in this section.

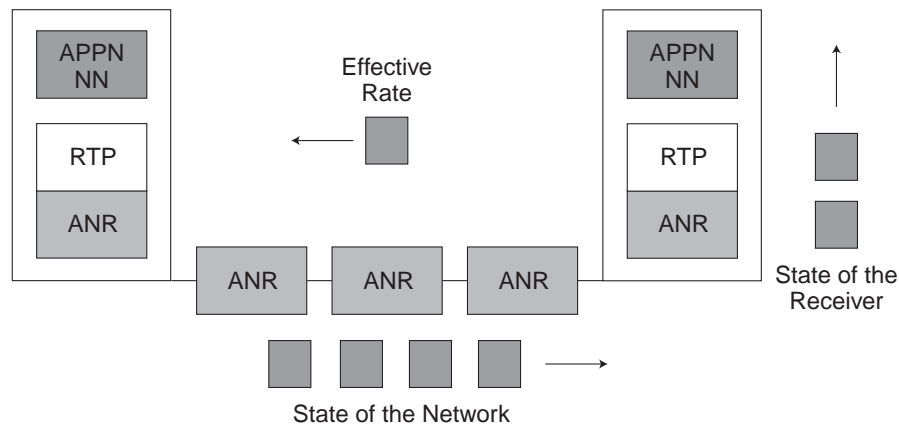
Memory and CPU Processing Requirements

First of all, a large amount of memory is needed to maintain the topology database and COS tables. As the size of the network grows, the number of potential paths increases significantly. This condition leads to a corresponding increase in memory usage, which in turn puts an upper limit on the size of the network. Although there are no firm measurements, the absolute upper limit is in the 200-300 NN range. This estimate is based on the size of the NN processor and its memory, as well as other configuration factors such as CP-to-CP sessions and configured cache sizes. Link state algorithm may cause a great deal of topology database updates (TDUs) in large and unreliable APPN networks.

Flow Control Issues

It was also well known that APPN HPR created severe network congestion issues because of problems associated with the original ARB flow control algorithm, ARB-1. The ARB-1 algorithm was ineffective when competing for bandwidth running over multiprotocol IP networks. ARB-1 was designed to be a proactive congestion control algorithm, which monitors round trip times and applies a smoothing algorithm to maintain flow control based on rate of the receiver, the objective being the avoidance of congestion and packet loss in the network (see Figure 3-11).


Figure 3-11 ARB Flow Control



The major problems with deploying ARB-1 in EE (HPR over IP) networks is that ARB-1 was designed to be a very conservative algorithm, which expected very little deviation in round trip times from averaged round trip times (otherwise known as smooth round trip time). Any deviations in round trip times from smooth round trip times were viewed as being extremely serious by ARB-1 and resulted in flow control rates being cut by 50 percent (sometimes more) for each deviation that occurred. This type of algorithm worked well in single-protocol native SNA APPN environments with extremely low error rates on links, but it was unacceptable and intolerant of sharing links with other protocols such as IP.

The following list summarize some of the major limitations with ARB-1 flow control:

- *Slow rampup*—ARB-1 was designed to ramp up slowly to ensure network stability. This results in bad performance for short connections where the data transfer is complete before ARB has a chance to ramp up to the available network capacity.
- *Lack of fairness*—ARB-1 was designed to be fair to all connections regardless of the rate at which they are sending. However, the algorithms used to convert an allowed send rate to a burst size and burst time can cause a lack of fairness. Simulation studies confirm that a connection running with a large burst size can use enough network resources to prevent a connection using a smaller burst size from ever ramping up to a fair share.
- *Poor performance on high-speed links (T3 and above)*—ARB-1 requires progressively better clock resolution to operate effectively at higher-speed links. At speeds higher than Token Ring, a clock resolution of better than 10 ms may be required to accurately determine the proper operating range. As a result of the loss of ARB control, implementations may not make effective use network resources at those speeds.
- *Poor fairness over short term with large number of connections*—Related to the design for slow rampup, it may take even longer for a connection to ramp up to its fair share if the network is already highly utilized by other connections.
- *Overreaction to losses*—The design assumption was that ARB-1 would reduce the send rate before congestion occurred; therefore, any packet loss was considered a sign of severe congestion. ARB-1 reacts severely to packet loss. By contrast, TCP uses packet loss as a normal indication of congestion. TCP also reduces its send rate in reaction to packet loss, but it recovers much more quickly. Tests of HPR and TCP sharing highly utilized network resources show the HPR traffic gets squeezed down to a minimal share of the network while TCP traffic gets almost all. This becomes an increasingly significant issue as customers build more multiprotocol networks that use the same resources for both TCP and SNA traffic.



Responsive Mode ARB (ARB-2), introduced in 1998 by the AIW, provides dramatically better performance and stability in HPR over IP (EE) networks. ARB-2 flow control is significantly more efficient than ARB-1 in environments that involve a multiprotocol IP network where different types of traffic using different flow control algorithms share some of the same physical bandwidth resources. More efficient rampup time for short connections, fairness to all connections, and minimized reaction to frequent packet loss can be obtained by having ARB-2 support enabled between RTP endpoints. ARB-2 improves data flow and allows HPR the ability to better compete with IP for bandwidth.

The ARB-2 enhancement is included and supported by the SNASw EE feature and is also supported in CS/390 V2R6 (with APAR OW36113 applied) and higher.

It is important to understand that ARB flow control levels are negotiated between RTP connection endpoints during RTP connection setup. Base ARB-1 support will be used unless *both* sides of the RTP connection support ARB-2. If you plan to have HPR-enabled APPN ENs downstream from an SNASw router that only support the original ARB-1 algorithm (for example, IBM AS/400 HPR support), then the RTP connection endpoints will be between the EN/HPR ARB-1 node (the AS/400 in this example) and the upstream EE-enabled NN server host RTP endpoint. In this example, ARB-1 (not ARB-2) will be the HPR flow control algorithm enabled between the AS/400 RTP endpoint and EE-enabled upstream NN server RTP endpoint.

It is highly recommended that HPR support be disabled on ENs downstream from an SNASw router in situations where an SNASw router is enabled for EE SNA transport over IP to an EE-enabled NN server upstream. That will allow ARB-2 to be the supported flow control algorithm type between the SNASw EE RTP endpoint in the router and the upstream EE host enabled for ARB-2 (z/OS CS control: HPRARB=BASE|RESPMODE).

Scalability Issues

The purpose of BrNN architecture and SNASw BX design stems from the fact that APPN NNs implement and maintain full APPN network topology database and full topology awareness of the entire network. NNs maintain information on the status of all links in the network and participate in searches through the network for resources.

Looking at this from the perspective of a remote branch network, however, there is really no need to know about the topology of the entire APPN network; any resource that is not located in the local branch itself should be located through the central site NN server if it is actually an available SNA resource somewhere in the network.

If the branch network contains no downstream APPN devices (that is, contains only dependent PU 2.0 SNA devices), then the node that connects to the central site really only needs to be enabled as a BrNN (SNASw BX) configured to serve dependent SNA devices using DLUR (DLUS is implemented on the APPN NN server host).

SNASw BX BrNN support allows the implementation of a node type that does not have to maintain complete topology awareness. It only has to maintain topology awareness for the downstream network below the SNASw BX router itself, and it does not need to participate in extensive network searches for resources. This allows a much simpler APPN implementation and a reduction in the total use of WAN bandwidth. This also allows an APPN branch implementation without the cost and overhead of having to implement full APPN NN routing functionality in the branch. Because SNASw BrNNs do not participate in network broadcast searches or topology updates (the SNASw BX router registers as an EN to the upstream NN server), the use of BrNN architecture allows for a much more scalable APPN deployment. Network reliability and stability is increased because BrNNs are immune from locate searches and broadcast storms associated with network broadcast

searches, which can result from repeated searches for nonexistent resources in large SNA networks. Broadcast storms and extensive locate searches can especially lead to network stability problems when these searches are repeatedly sent over limited-capacity WAN links.

Another advantage of the BX architecture is that the BrNN only needs to have a single upstream CP-to-CP session to an upstream NN server. The SNASw BX BrNN must explicitly define the links to other NNs serving as backup NN servers, but all other ENs upstream from SNASw should use connection network to support dynamic link connections between ENs (as was covered earlier in this chapter). The SNASw DLUR function can either dynamically connect to alternate upstream NN/DLUS servers serving as backups, or can predefine the connection to a single alternate NN server acting as a backup DLUS. The purpose of this design approach is so that the BrNN has only a single link in the upstream network over which it can send search requests for unknown resources to the NN server; the NN server acts like a default router. Thus the BrNN need only maintain a topology database for the downstream branch network, and it simply forwards requests for any other resources not local to the SNASw BX router over the CP-to-CP session link to its upstream NN server.

Looking further at BX architecture we see that it really has a “dual personality.” To the NN enterprise server (CS/390, for example), BX looks like an EN. This means that the only other NN seen by an enterprise NN server would be another enterprise NN server. Topology broadcast traffic would therefore be limited to those NN servers and would not be sent to SNASw BX. SNASw BX, however, looks like a NN to the nodes and SNA devices downstream. This means that BX sees no other NNs downstream and, therefore, sends no topology broadcast traffic.


BX works with any downstream EN (except a CS/390 VTAM or PBN host—see the “SNASw General Considerations” section) to select the session path from the end user to BX. BX works with the CS/390 NN server to select a session path between BX and the application host. Links to nodes downstream of a BX are called *downlinks*. Likewise, links to upstream devices in the host network are called *uplinks*. SNASw uses a unique method to distinguish uplinks from downlinks. Every link that is defined to SNASw is automatically considered an uplink. Links that are dynamically generated by SNASw as a result of an SNA device (EN, LEN, or PU 2.0) initiating a connection into SNASw are automatically considered downlinks.

Network Design Considerations for APPN NN to SNASw BX/DLUR Migration

This section reviews several network design issues that should be considered by organizations currently running large APPN NN networks and migrating to Cisco SNASw. The key point to note here is that adding Cisco routers with SNASw function requires a careful insertion strategy if there are currently a number of cascaded NNs in the network path from downstream SNA devices to the hosts.

The SNASw BX router must not have traditional APPN NNs connected below it in the network topology. As shown in Figure 3-12, you should replace APPN NN with Cisco SNASw routers at the bottom-most layer of the network whenever possible (depending on whether an existing DLSw+ network is in place). If further removal of network-based NNs is desired, the next step is to design for direct links between the SNASw BX nodes and the CS/390 NN server hosts instead of having a complex cascaded NN topology in the middle, because all other links to application EN hosts should be enabled dynamically using SNASw connection network support as previously discussed in this chapter.

Figure 3-12 BX Network Design for Migration



To use SNASw in network designs that include an existing DLSw+ network for SNA transport over the WAN, SNASw functionality can be deployed either in the same data center hub-end routers being used as DLSw+ peer aggregation points (taking into account sizing and scaling considerations for the additional overhead of running SNASw and DLSw+ in the same routers) or in separate routers. Running SNASw and DLSw+ in the same router allows APPN COS to IP ToS mapping to be supported for outbound connections over DLSw+.

For enterprises with multiple data centers where traffic is consistently routed between the data centers, you should consider extending SNASw BX and EE functionality to regional offices and aggregation points.

SNASw could also be deployed in separate routers than the routers supporting existing DLSw+ functions. Running SNASw and DLSw+ in separate routers may be a good approach for large networks that have significant amounts of multiprotocol traffic. The approach might be more beneficial for change management control, network availability, and avoiding single points of failure. However, running SNASw in separate routers from central site DLSw+ routers would necessitate another layer of SRB paths to support LLC traffic between the DLSw+ and SNASw routers.

Some additional APPN NN-to-SNASw migration considerations are the following:

- Target application hosts should be defined as ENs.
- SNASw links to NN server hosts (for CP-to-CP session support) should be explicitly configured in SNASw.
- SNASw connection network support (previously covered in this chapter) should be used to support dynamic links to all other application EN hosts.

- If you plan to have EE (HPR-only) connections adjacent to interchange transmission groups or if any of your hosts are connected to ICNs, there exists an IBM APAR for OS/390 and CS/390 releases prior to V2R10 (and releases prior to IBM APAR OW44611) that did not allow sessions to cross-domain subarea partners to exit an ICN via an HPR connection unless the connection was only one hop away from the target EN.
- Links from SNASw to ICNs should be explicitly predefined even when running IBM OS/390 and CS/390 V2R10 or higher (or when IBM APAR OW44611 is applied) because there is one scenario that will not support HPR on an APPN link immediately adjacent to an interchange transmission group. This is when an ICN defines a connection to a connection network (VRN) and a session is attempted from the subarea network through the ICN and then into APPN over the VRN. Refer to IBM APAR OW44611 for more information regarding this limitation.

If separate channel-attached CIP or CPA routers are defined, traffic can be bridged at Layer 2 from SNASw to the enterprise server across the channel-attached CIP and CPA routers. As mentioned before, SNASw BX and other application EN target hosts can be defined as nodes in a connection network.

Leveraging SNASw EE for APPN NN-to-SNASw Migration

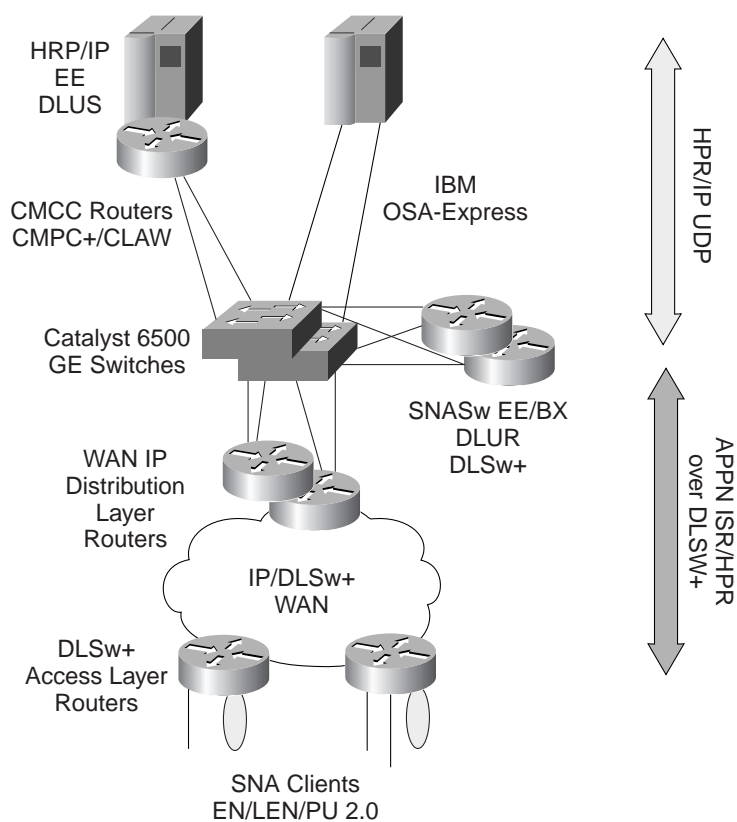
EE was created within the open framework of the AIW and submitted to the IETF as RFC 2353 in May 1998, and it has been commercially available since 1999. Cisco SNASw implements the EE feature as well as IBM CS/390 V2R6 (with APAR OW36113 applied) and higher. EE is a logical connection that represents IP connectivity from an EE-enabled IBM S/390 or zSeries enterprise server to the SNASw EE router end to end. EE allows the enablement of IP applications and convergence over a single network transport (IP) while preserving customer SNA application and SNA client endpoint investments.

From SNA's perspective, EE is just another DLC type. From IP's perspective, EE is just another UDP connectionless transport layer application operating at Layer 4 of the OSI model!

EE is an open technology that totally integrates SNA devices into an IP network. Consolidating parallel networks onto a single standard IP network provides a significant savings in equipment and administrative costs. With EE, IP can be extended all the way from the enterprise server to SNASw EE-enabled routers in remote branch offices, distribution layer routers, or data center routers. Implementing EE leverages the fault tolerance, rerouting, and high-performance capabilities of an IP network and IP running in the enterprise server while greatly simplifying management. At the same time it ideally positions the enterprise for adoption of emerging IP multiservice technologies such as Cisco AVVID (for example, voice over IP [VoIP]) and integration of high-speed data center and campus IP Layer 3 support.

With EE transport, we generally expect to see two basic models of network designs. In the first model (see Figure 3-13), EE is used in conjunction with an existing DLSw+ network to remote branches.

Figure 3-13 EE Migration Model 1: DLSw+ to the Branch



Because both EX and DLSw+ provide options to transport SNA over IP, one might ask why you would want to combine both into a single network design. There are some very good reasons for doing this. In Figure 3-13, SNASw is running in addition to DLSw+ in the data center routers to provide necessary SNA session routing, while the SNASw EE feature is enabling the transport of SNA traffic over IP natively from the SNASw router to the enterprise server. For the existing Cisco DLSw+ customers, this combined SNASw/DLSw+ solution approach provides the ability to leave the existing remote DLSw+ router software unchanged while only enabling data center (or aggregation layer) router software to support both SNASw and DLSw+ functions.

SNASw support for BrNN continues to provide emulated NN services for all downstream EN/LEN nodes out in the remote branch network, while SNASw DLUR provides support for PU 2.0 devices downstream. By maintaining DLSw+ outbound, the organization continues to leverage the value of its consolidated SNA and IP network, while it begins to migrate safely and cost-effectively to a full IP infrastructure. Using IP upstream simplifies design configuration and provides the opportunity for integration of the enterprise network with Token Ring-to-Ethernet LAN migration within the campus layer infrastructure.

HPR over IP requires only a single link definition in CS/390. Also, IP, rather than SNA, handles session path switching, providing faster session rerouting around link failures. SNASw rerouting enables a highly flexible and scalable data center design where there are no single points of failures between data center SNASw EE routers and EE-enabled CS/390 enterprise server.

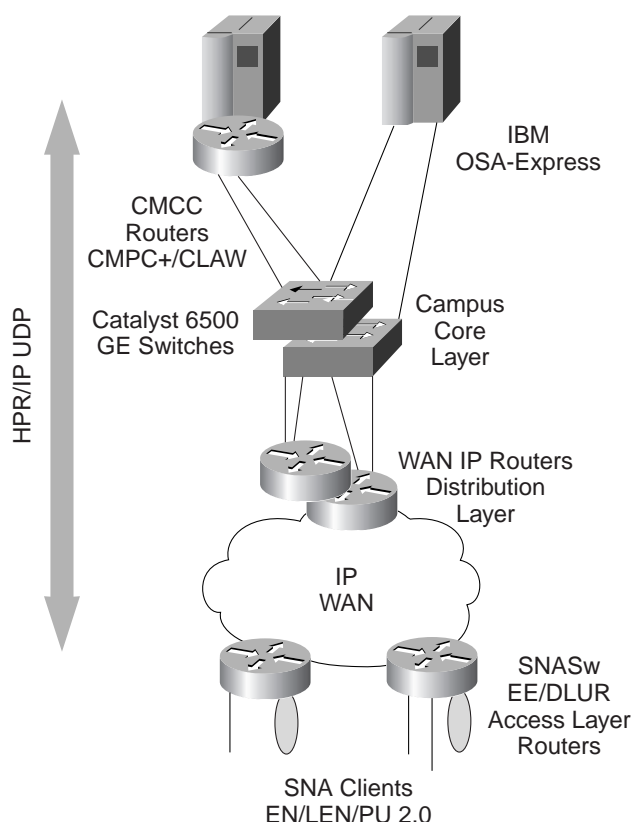
This design option illustrates the use of both Cisco channel-attached routers (CIP/CPA) for the IP uplink to the mainframe server and IBM's Gigabit Ethernet OSA-Express. Both solutions provide viable design alternatives. Typically, when the IBM S/390 (or zSeries) Parallel Sysplex is running OSA-Express, a Cisco CIP or CPA will not be involved and a Cisco Catalyst 6500 Series Gigabit Ethernet switch will directly attach to the OSA-Express network interface card (NIC) in the enterprise server (S/390 G5, G6, or zSeries mainframe). In another scenario, if a CIP or CPA had previously been installed to replace a FEP, the OSA-Express could handle all IP traffic and the CIP could continue to handle all remaining SNA over LLC traffic (the CIP and CPA can also continue providing support for IP transport to the host).

It is important to note here that SNASw EE supports SNA COS to IP ToS mapping (SNA transmission priority to IP precedence mapping) for both inbound and outbound traffic in a bidirectional fashion between the S/390 EE-enabled host and the SNASw EE router. If you are using DLSw+ for downstream SNA WAN transport between a combined SNASw/DLSw+ router and remote SNA devices, COS/ToS is only mapped in an outbound direction (from the SNASw/DLSw+ router outbound). On the reverse path upstream to the host from the remote DLSw+ router supporting SNA device connections, COS/ToS mapping cannot occur.

Finally, there are some basic considerations to this design approach: IBM mainframes may require an operating system upgrade for this design to work (CS/390 V2R6 with APAR OW36113 or higher is required for EE); APPN/HPR must be running in CS/390 on the mainframe, and there must be IP support enabled on the hosts.

The second model, shown in Figure 3-14, demonstrates how SNASw EE can be extended from the EE-enabled enterprise server all the way to the SNASw EE router in the remote branch office. Across the WAN, the SNA traffic is transported using dynamic IP routing protocols such as Open Shortest Path First (OSPF). The end-to-end HPR flow, error, and segmentation control would be performed between the SNASw EE router at the branch and the enterprise NN EE host server.

Figure 3-14 EE Migration Model 2: EE to the Branch



The network design scenario shown in Figure 3-14 optimizes SNA response time and availability, because EE is the preferred technique for transporting SNA over the IP backbone end to end. IBM S/390 G5, G6, or zSeries mainframes are already installed, with Cisco Catalyst 6500 Gigabit Ethernet switches connected to IBM OSA-Express also in the plan. The OS/390 software is already at a level that supports EE in the enterprise server. The result is an IP transport network from the branch all the way to the S/390, with nondisruptive rerouting for network outages.

In addition, the network is enabled to differentiate QoS services and prioritize within SNA traffic (interactive SNA versus batch), as well as between SNA and IP traffic. Unlike the combined SNASw/DLSw+ approach, this differentiation occurs in a bidirectional fashion between the EE-enabled enterprise host and SNASw EE router (at the remote branch). Differentiated services allows the service policy agent within CS/390 to enforce QoS policies based on time of day, day of week, end user, and so on, providing more flexibility than traditional SNA COS support can.

By using EE you now have a full IP network from the remote branch supporting SNA client directly into the host. With SNASw EE in the branch and EE running in the enterprise server, HPR over IP is supported across the entire network. SNA traffic is carried natively over IP and does not need DLSw+ for SNA transport over the WAN (DLSw+ is completely unnecessary in this design).

The key advantage to this optimal design approach for SNA transport is the creation of a full IP network infrastructure from the SNA client to the SNA application running on the target application host. With IP running end to end, there is logically no real single point of failure anywhere in the network except the SNASw branch router itself (a design could incorporate SNASw branch router redundancy using Hot Standby Router Protocol [HSRP] support for multiple SNASw routers on a remote LAN segment).

Customers Migrating from Subarea SNA (FEP) to SNASw

A FEP is a class of communications controller (PU 4) that off-loads host mainframes from many communications functions. Traditionally, SNA routing has been done on FEPs, but the majority of organizations are choosing to migrate from FEPs to higher-speed, multiprotocol routers to save money and to position their networks for S/390 IP applications. The SNASw solution allows migration from a FEP-based data center supporting both independent LU (LU 6.2) traffic as well as traditional subarea SNA traffic to a consolidated data center that supports SNA and TCP/IP applications concurrently.

As was discussed in an earlier section of this design guide, one of the first things you want to do before considering implementing Cisco SNASw is to evaluate the SNA routing requirements in your network. If SNA routing is required to route client session requests directly to target application host LPARs, and FEPs are currently installed, you must make a decision to either keep some FEPs in place (for SNA routing) or replace FEPs with Cisco CIP- or CPA-attached routers (or IBM OSA-Express) and replace the SNA session and application routing functionality provided by the FEPs (and dependent SNA boundary function for PU 2.0 nodes) with SNASw.

Because FEPs are costly to run and are not compatible with an optimal IP network design, replacing FEPs has become a very high priority in enterprise network migrations. If your network currently uses FEPs for native SNA routing, and you are migrating your data center from FEPs to CIP and CPA platforms, then you need SNASw somewhere in your network. As discussed previously, you can deploy SNASw all the way to the remote branch (as an alternative to DLSw+) or only where you currently have FEPs (in which case you can still use DLSw+ in your branch offices to transport SNA traffic to the data center over the WAN).

The Cisco CIP and CPA with SNASw cannot replace all FEP functions. The Cisco IOS Software and the CIP and CPA can address most of the key FEP functions while providing a higher-performing, multipurpose channel gateway. In replacing FEPs in the data center, a CIP or CPA can serve as the channel gateway between the SNASw routers and host mainframes.

In multihost environments the SNA routing support in SNASw allows you to minimize or eliminate your dependency on FEPs and NCP software while migrating to a Cisco CIP/CPA or IBM OSA-Express if you are migrating to Cisco Catalyst 6500 multilayer switches in the campus.

For functions not addressed by the Cisco CIP, CPA, and SNASw, one or more FEPs may still be required. You should take the following issues into consideration when replacing FEPs with Cisco CIP or CPA with SNASw:

- The Cisco CIP and CPA are not SNA PUs and do not provide subarea SNA boundary function in and of themselves. The solution here is to implement SNASw DLUR support in combination with the CIP or CPA. SNASw DLUR replaces SNA boundary function previously provided by FEPs for peripheral PU 2.0 SNA devices.
- SNI is an SNA-defined architecture that enables independent subarea networks to be interconnected through a gateway. SNI connections require a FEP in at least one connecting network. The Cisco IOS Software allows connection to an SNI gateway host but does not provide SNI gateway functionality. One solution for replacing SNI is to migrate to an APPN network interconnected by a border node (either extended or peripheral), which allows networks with different topology subnets (NETIDS) to establish CP-to-CP sessions with each other. SNASw does not play any role whatsoever in host-to-host border node implementations. Cisco routers and multilayer switches can provide IP transport between hosts that have implemented extended border node

(HPR/IP EE) support, or they can provide DLSw+ SNA WAN transport for bridged LLC traffic from non-HPR/IP (EE) border node connections between mainframe hosts (that is, hosts not implementing extended border node EE support).

- Although the Cisco IOS Software can duplicate some other special protocols supported by the FEP (such as asynchronous and bisynchronous tunneling), conversion to SNA is not provided. The CIP and CPA can provide TN3270 Server support, which can enable conversion from TN3270 client connections over IP to SNA upstream to the host. For more information, see the *TN3270 Server Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/tn3270sr/tech/tndg_toc.htm.

Subarea SNA to SNASw DLUR Migration Considerations

SNA DLU sessions (PU 2.0) can be supported using DLUR/DLUS function. Cisco supports older DLU traffic across APPN networks through the SNASw DLUR feature. SNASw DLUR provides the opportunity for customers to migrate from IBM FEPs boundary function support in subarea SNA environments today to boundary function support over Cisco powered networks using the DLUR capability of SNASw.

This enables DLU session traffic to take advantage of dynamic directory and topology functions in APPN. DLUR addresses the DLU migration problem by removing the logically adjacent restriction, allowing a DLU to be remotely attached to a subarea node and receive SSCP services. The SNASw DLUR feature, combined with DLUS services in the enterprise NN server CS/390 host, provides the benefits of SSCP services and allows the data and BIND path for the SNA LU-to-LU session to be different from the SSCP-to-PU and SSCP-to-LU control sessions. This path difference allows the LU-to-LU sessions to take full advantage of APPN route selection dynamics. By putting the DLUR function in the Cisco router, remote SNA controllers need not be upgraded to support APPN/DLUR. The DLUs continue to remain visible to NetView for network management because CS/390 maintains awareness through DLUS-to-DLUR session flows.

All DLUs, and the PUs that support them, require sessions to their owning SSCP. These sessions carry various control and management requests such as INIT-SELF, NOTIFY, NMVT, and USS messages. They always take the form of SSCP-to-PU and SSCP-to-LU sessions, which, before SNASw DLUR support, flowed entirely within a single CS/390's SSCP subarea domain. That meant that a PU serving DLUs always had to be directly connected either to its owning SSCP domain or to a FEP NCP owned by that domain. Cross-domain or cross-border SSCP ownership of DLUs was totally out of the question.

Another restriction affecting DLUs is that routing in a subarea network is always done at the *subarea level*. In other words, any session involving a DLU must pass through the same adjacent subarea node as the SSCP-to-LU session, even if the DLU happens to reside in an APPN node!

SNASw DLUR support eliminates both of these restrictions by providing a number of SNA architecture functional enhancements. With SNASw DLUR, sessions between each DLU (or PU) and its SSCP are now encapsulated within an LU 6.2 pipe consisting of a pair of sessions between the CPs in the DLUR and DLUS nodes (using the mode name CPSVRMGR and the APPN COS SNASVCMG). The DLUR/DLUS pipe can carry a number of SSCP-to-PU and SSCP-to-LU sessions and does not need to be between adjacent CPs. The pipe can be carried on a base APPN ISR or HPR ANR connection and, unlike subarea SNA, can cross APPN network boundaries.

With SNASw DLUR, LU-to-LU session routing is now performed wholly by the APPN function and does not require the subarea boundary function to be at an adjacent subarea node. In fact, the SNASw DLUR router itself provides the boundary function. When a primary LU requests a search for a DLU, it normally receives a positive response from the DLUS, not the DLUR. The response indicates the DLUS as being the NN server for

the DLU. The route is then calculated directly to the DLUR by the NN server of the primary LU. In some cases (where the DLUR supports cross-network DLUR/DLUS control sessions) the DLUR itself may respond to a search, in which case the CP name and NN server name given are the correct ones.

Note: Because the DLUS presents itself as the NN server for the DLUs, it must always be the NN directly connected to the downstream SNASw DLUR router.

SNASw DLUR support requires no changes to the existing host applications or DLU terminals in the network.

One major restriction that exists in subarea SNA networks is the limitation of only supporting a maximum of 64,000 network addressable units (LUs, PUs, and CPs) within a single SSCP domain. This limitation was addressed by the enhanced network addressing APPN support provided in CS/390 V2R5 (with APAR OW32075) and higher, which allows SNASw DLUR- and CS/390 DLUS-served LUs above the 64,000 network addressable unit address limitation in subarea SNA networks.

SNASw DLUR Design Considerations

The SNASw DLUR function greatly improves the flexibility available to the network designer by offering new options for both routing and connectivity. However, there are several points that must be taken into consideration before implementation:


- The SNASw DLUR connection to the upstream DLUS must be established over an APPN network with no subarea hops in between.
- LEN connections are not permitted over a DLUR/DLUS pipe.
- The primary LU CP and its NN server must support the session services extensions APPN option set. Only CS/390 supports session services extensions; thus functions such as the AS/400 primary LU support cannot be used with DLUR LUs unless the AS/400 is in the subarea network and therefore uses a VTAM ICN as its APPN primary LU node.

Customers Migrating from Token Ring to High-Speed Ethernet Campus LAN

Long ago, Ethernet eclipsed Token Ring as the dominant enterprise LAN technology for new installations. It did so because it was less expensive and was offered by a larger number of vendors than Token Ring. However, some of the largest enterprises in the world continued to maintain and enhance their installed Token Ring networks, at least within a portion of their networks. Often their reasoning was driven by the fact that, historically, Token Ring was the preferred technology to support mission-critical SNA traffic. Token Ring is very stable and it scales gracefully to support a large number of users.

Token Ring has now become outdated and is a niche technology. The number of vendors providing Token Ring solutions is shrinking, and for some products there is only a single vendor still in the market. The prices for Token Ring solutions remain high compared to Ethernet-based solutions and, because of the lack of competition, will continue to remain so. Finally, Ethernet is better equipped to support emerging networking applications and technologies such as gigabit speeds, multimedia, multicast, and voice/data integration applications. Therefore, the majority of enterprises with Token Ring installed are implementing plans now to migrate to Ethernet as quickly as possible.

A popular misconception of some proponents for maintaining a Token Ring infrastructure is the belief that sticking with Token Ring is a “zero-investment” decision. In reality, the decision to remain with Token Ring implies a continued investment in the technology, with the purchase of new Token Ring NICs for each new PC workstation installed, in addition to the purchase of new Token Ring switches and routers to support the demand for increased network bandwidth. The financial metrics of this dictate a migration to Ethernet over time.



The migration can be swift or slow, depending on the particular needs and priorities of the organization. The Token Ring migration is often coupled with other infrastructure changes, such as the elimination of legacy protocols in the network (in favor of SNA over IP using SNASw EE transport) or the refresh of desktop PCs.

One of the primary reasons for the prolonged dominance of Ethernet technologies is the price/performance curve available to users of Ethernet-based products. In the early days of shared Ethernet environments, enterprises could point to the higher speeds and more deterministic operation of Token Ring as justification for continuing to invest in the technology. However, with the more open nature of Ethernet and the continued investment by a number of vendors, the price of Ethernet-based solutions declined almost exponentially while the transmission speed continued to increase. Shared 10-Mbps bandwidth was quickly replaced by switched 10-Mbps bandwidth. Switched 100-Mbps is now becoming commonplace at the desktop, and Gigabit Ethernet solutions are being implemented using Catalyst 6500 multilayer switches on the campus backbone. Token Ring simply has not kept up. The price of 16-Mbps Token Ring switch ports is still higher than 100-Mbps Ethernet switch ports. High Speed Token Ring (HSTR), a proposed standard for 100-Mbps Token Ring, has not taken off and there is no talk of a gigabit Token Ring solution. The price/performance curve will continue to favor Ethernet solutions going forward.

The message from industry analysts and organizations that have undergone the migration is clear—make the decision to migrate right now! To delay the decision means risking your ability to react in the future to deploy new networking applications and new networking technologies. In the long run, delaying the migration will cost more and will place your network at higher risk than if you get started today.

For a more detailed discussion and business case, see *Token Ring-to-Ethernet Migration* at www.cisco.com/warp/public/cc/so/neso/ibso/ecampus/trh_bc.htm.

How SNASw EE Can Leverage Token Ring-to-Ethernet Migration

Token Ring and SRB have traditionally been utilized to maintain multiple concurrently active redundant paths (RIFs) in a bridged network to the mainframe without the inherent problems with routing loops that exist with LLC2 Layer 2 bridged transport to the host over Ethernet (for more detail, see *Ethernet DLSw+ Redundancy* at www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/prodlit/appxc_rg.htm). The foremost advantage that SNASw EE provides for customers migrating from Token Ring to Ethernet is the ability to consolidate SNA onto a single IP network transport model. This eliminates the need for using LLC2 source-route bridged Token Ring transport of SNA packets from the network to mainframe enterprise servers.

SNASw EE leverages the inherent availability features of IP at Layer 3 (versus source-route bridged LLC traffic over Token Ring at Layer 2) to provide failure-resistant SNA application access *regardless* of the underlying LAN or WAN medias. When multiple paths to enterprise servers exist, normal IP reroute capabilities maintain all connections and dynamically switch SNA client sessions to the application host directly without session disruption, thus avoiding all single points of failure in the network.

High-speed IP data center and Token Ring-to-Ethernet migration is leveraged by SNASw EE using a number of different technologies for IP Layer 3 transport of SNA traffic to the host:

- Cisco Catalyst 6500 multilayer switches and IBM OSA-Express
- Cisco CIP
- Cisco CPA

If an enterprise needs to transmit multimedia data in the gigabit range, IBM OSA-Express and the Cisco Catalyst 6500 switch can work together with IBM G5, G6, and ZSeries series hosts to transmit high-speed IP data. With EE running in these host platforms, the Catalyst 6500 can support Gigabit Ethernet speeds and provides the best choice for connectivity to IBM's OSA-Express NIC in the enterprise server.

The IBM OSA-Express adapter is going to become the method of choice for attaching a S/390 (G5 or G6) and zSeries host to a TCP/IP network. The new architecture eliminates the limitations of the channel protocols and puts the S/390 on the same plain as large UNIX servers. By using Queued Direct Input Output (QDIO), the Gigabit Ethernet and Fast Ethernet cards have direct access to the 333-MBps CPU buses. This is considerably faster than the current ESCON technology, which relies entirely on channel protocol support.

By rewriting the TCP/IP stack to use direct memory access (DMA) against the OSA-Express buffers, IBM has eliminated many of the previous performance issues associated with buffer copies. This results in better throughput, at reduced CPU resource consumption.

IBM has also reduced the amount of configuration that is required for TCP/IP passthrough, loading the parameters from the TCP/IP profiles dataset. This eliminates the need to use the OS/2 or Windows-based OSA/Support Facility (SF).

The OSA-Express supports the Service Policy Server in S/390. The OSA-Express has four output queues. Each queue is associated with ToS. Application data is prioritized by the Service Policy Server, and data is queued in priority. ToS and Diffserv bits are set and read by the Cisco network, providing end-to-end QoS.

Customers should consider the use of the OSA-Express for Gigabit Ethernet TCP/IP connectivity to the mainframe. It provides higher throughput than is possible with the Cisco CIP or CPA. The types of TCP/IP access that should be considered include high-volume FTP, APPN/HPR over IP (SNASw EE), and access to a mainframe-based TN3270 Server. Throughput can be expected to be three to four times greater for bulk data transfer, at a reduced CPU consumption of up to 25 percent (interactive data flows will not see the same level of improvement). TN3270 transactions will run up to 10 percent faster, depending on message size, with a reduced CPU consumption of up to 10 percent.

High-speed IP data transport (up to Fast Ethernet 10/100 Mbps wire speed) is also supported by the Cisco CIP and CPA. When a host mainframe is configured with EE, the CIP and CPA can also provide IP connectivity to the host for downstream SNASw routers running the EE feature.

Customers are much better off using a CIP or CPA in the following instances:

- *Non-IBM mainframes*—The CIP and CPA can be used with any mainframe that supports the Enterprise Systems Connection (ESCON) or bus and tag channel protocols. This is 100 percent of all IBM and plug-compatible boxes. The OSA-Express is not an option for non-IBM mainframes.
- *Older mainframes*—The CIP and CPA can be used on the approximately 60 percent of mainframes that do not support the OSA-Express.
- *Older operating system releases*—The CIP and CPA can be used with any currently supported operating system release.
- *Aggregation of TCP/IP and SNA traffic*—The CIP and CPA represent efficient usage of ESCON card cage resources. The interface cards in the router can be used to aggregate LAN and WAN traffic, and the combined traffic can be efficiently transported across the ESCON or bus and tag channel.
- *Few or no available ESCON card cages*—Because the CIP and CPA can be attached to an existing ESCON channel via an ESCON Director, no additional frame is needed for card cages.
- *Offload processing*—The dedicated CPU and memory of the CIP or CPA can be used to offload processing from both the router and the mainframe. The TN3270 Server application can be used to offload the protocol conversion duties from the mainframe. The TCP/IP Offload function can be used to offset the huge inefficiencies associated with the mainframe TCP/IP stack in older (V2R4 and earlier) CS/390 releases.

SNASw Migration Scenarios

Scenario 1—Large Branch Network Migration to SNASw EE

A large insurance company had recognized that migrating to an IP backbone and using TCP/IP as the primary transport protocol would lead to enormous productivity increases while shortening application development cycles and lowering costs. The company executives were aware of the alliance between Cisco and IBM and wanted to take advantage of the cooperative engineering teams to design a new network consisting of Cisco routers and IBM enterprise servers outfitted with OSA-Express Gigabit Ethernet LAN adapters.

Business Requirements

Any proposed network design had to incorporate a number of requirements to accomplish the business goals for this large insurance company:

- Support more than 2000 remote branch sites
- Handle SNA traffic between AS/400 systems at each branch and the main data center
- Allow for efficient connectivity for branch-to-branch SNA and IP traffic
- Permit easy rollout of new TCP/IP applications without impacting performance and throughput for current SNA-based applications
- Allow for a methodical, planned replacement of all IBM routers with Cisco network platforms while maintaining connectivity during the migration period
- Design for management from both an SNA (data center operator) perspective as well as IP (network operations) perspective

Network Analysis

To create the new network design, the engineers used a series of questions to help analyze the current network environment:

- *Is APPN required?*—APPN (and APPN/HPR) was already in place throughout the network. Because one of the major goals of the new network was to migrate to an all-IP backbone, replacing APPN routing with an SNA-over-IP solution was paramount.
- *How do we design for maximum availability and scalability?*—The network had to support more than 2000 branch locations. The current network used APPN and OSPF routing. It was decided to use OSPF routing in the new design and supplement the network routing by running OSPF on the data center hosts to facilitate establishing alternate routing paths to the data hosts dynamically. This meant that a hierarchical OSPF design with core, distribution, and access layers would best provide the performance, reliability, and availability requirements for this customer.

- *How should we connect the Cisco data center routers to the IBM hosts?*—The customer, IBM, and Cisco jointly determined that the best solution was to design a redundant, switching and routing network front-ending the data center hosts. The data center hosts would form a sysplex environment with multiple LPARs and interconnected CPUs. Access to the hosts would be via OSA-Express Gigabit Ethernet interfaces attached to Cisco Catalyst 6500 Layer 3 series switches outfitted with MSFCs. This combination would handle both SNA-based application traffic and TCP/IP-based application traffic all over a single IP backbone. Because the hosts are already APPN NNs and running OS/390 V2R8, it was decided that an all-IP network could be implemented using EE.
- *What should we use for transporting SNA traffic from remote sites to the data center?*—EE requires that HPR/IP be used as the transport mechanism. SNA application data is transported in IP/UDP frames. An additional benefit is the host automatically sets the IP precedence ToS bits according to the COS table entries defined for the session requests. This melding of the SNA COS prioritization scheme with the IP network prioritization scheme provides the best opportunity for maintaining current response time requirements while allowing the company to move forward with its IP application and voice, data, video implementation plans.
- *Where should we place SNA features such as DLSw+ and SNASw?*—As noted previously, the current network already implemented APPN routing within the network at Layer 2. There were also significant branch-to-branch SNA traffic requirements. This need, coupled with the desire for an all-IP Layer 3 network, led to a decision to implement SNASw at the branch access level. The IP backbone network would form an SNASw connection network for APPN traffic. The data center host NN server would make SNA routing decisions, and the actual data path for SNA application session traffic would be established over dynamically built links over the IP connection network.
- *How will we manage the integrated network?*—Because the customer had data center operators who were familiar with mainframe-based Tivoli NetView for OS/390 and wanted visibility into the IP router network, Cisco Internetwork Status Monitor (ISM) would be installed and configured to use SNMP to extract and display status information in a familiar form for these users. CiscoWorks2000 Routed WAN Bundle would also be installed on a UNIX server. This software would be used by the network operations staff for inventory management, configuration management, Cisco IOS Software management, and network troubleshooting.

Design Rationale

Based on the objectives of this large organization and an analysis of its current network design, Cisco engineers derived the following general points that would drive this new network design process:

- A hierarchical design would be stable, scalable, and manageable
- Parallel host connectivity would create an easy migration
- The plan must permit the redesign of an old IP addressing scheme to support new design
- The plan must permit redesign of a Frame Relay network for easy migration
- A converged IP network would support current SNA traffic and allow rapid deployment of new IP-based applications

SNASw played a large role in this network migration design because of its many features: BX, EE, HPR, automatic COS-to-IP precedence mapping, Layer 3 IP end-to-end routing, and Layer 4 HPR for reliability. Scalability of this network design was of primary importance. Therefore, SNASw was used because of its EN/NN emulation and its effect in an IP network:

- NNs run only in VTAM CMCs
- Topology updates and broadcast locates are eliminated from the WAN
- IP connection network attachment
- Dynamic link setup using connection network


The EE feature of SNASw would provide SNA-over-IP transport throughout the design because of its reliability, ability to support nondisruptive rerouting of SNA sessions around failures, and ability to support branch-to-branch traffic routed through region.

The Migration

The migration involved two phases. Phase 1 was to IP-enable the data center. Phase 2 involved the network migration in the branches in five steps:


- Step 1. Install data center WAN routers (see Figure 4-1)
- Step 2. Install regional backbone routers (see Figure 4-2)
- Step 3. Install regional distribution routers (see Figure 4-3)
- Step 4. Convert the branches and then the regions (see Figure 4-4)
- Step 5. Decommission the region IBM 2216 routers (see Figure 4-5)

Figure 4-1 Install Data Center WAN Routers



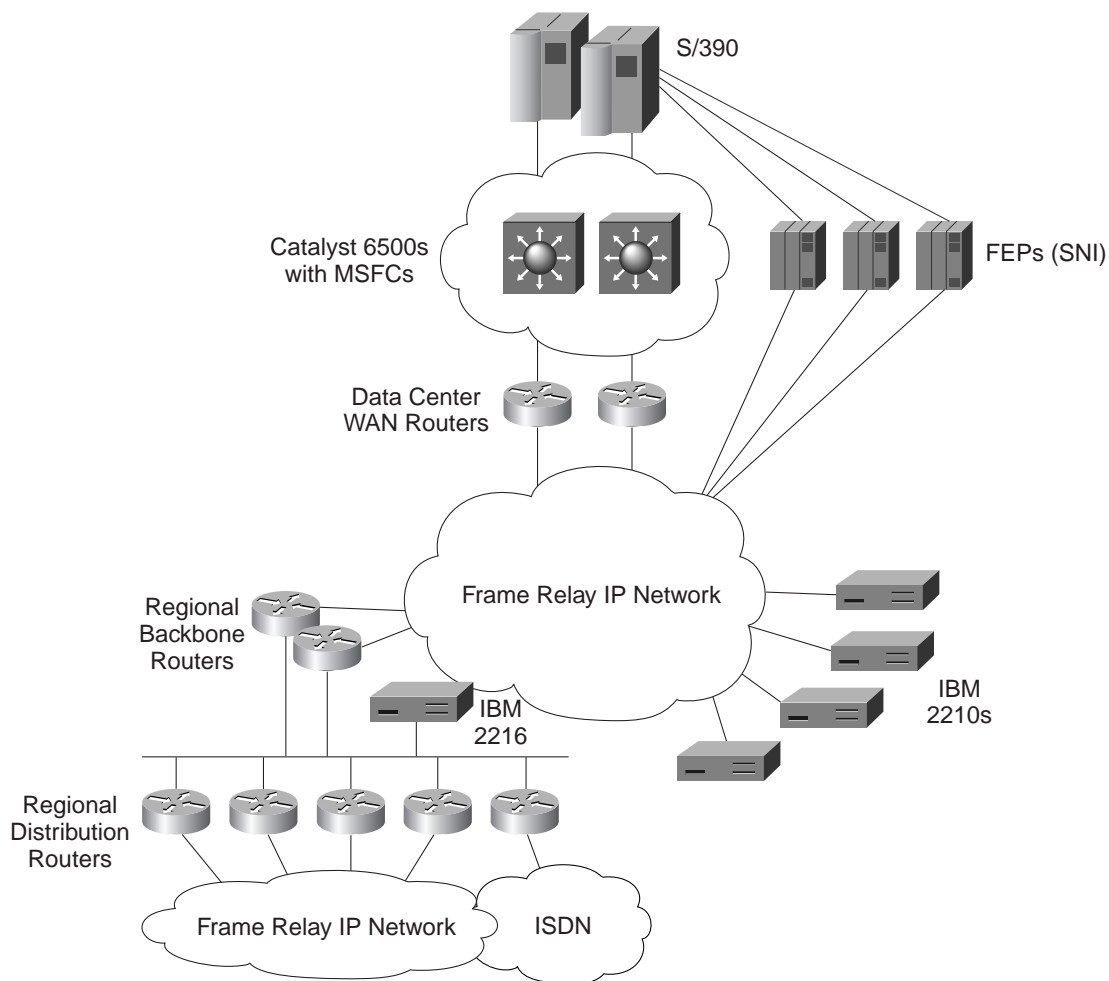
When the data center was prepared and the OSA-Express adapters were installed, the network migration could begin. The first step involved creating parallel network access to the sysplex using Cisco Catalyst 6500 switches. Cisco 7200 Series WAN backbone routers were installed and connected to newly provisioned circuits into the Frame Relay network.

Figure 4-2 Install Regional Backbone Router



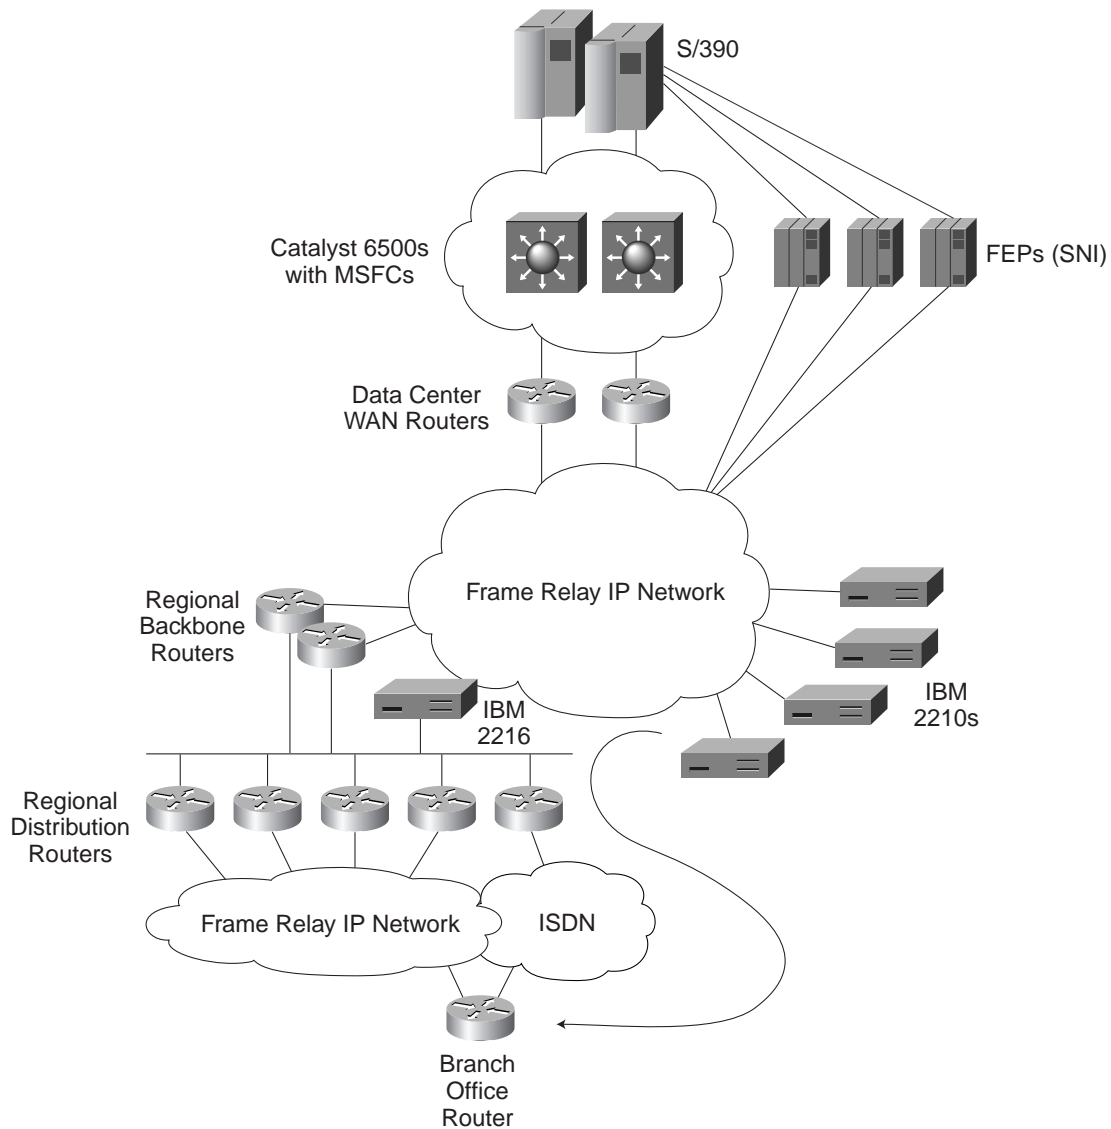
The next step was to install new equipment at each of the regional sites. First the regional backbone routers were added to the first of 10 regions. Connectivity was established back to the data center WAN backbone routers.

Figure 4-3 Install Regional Distribution Routers



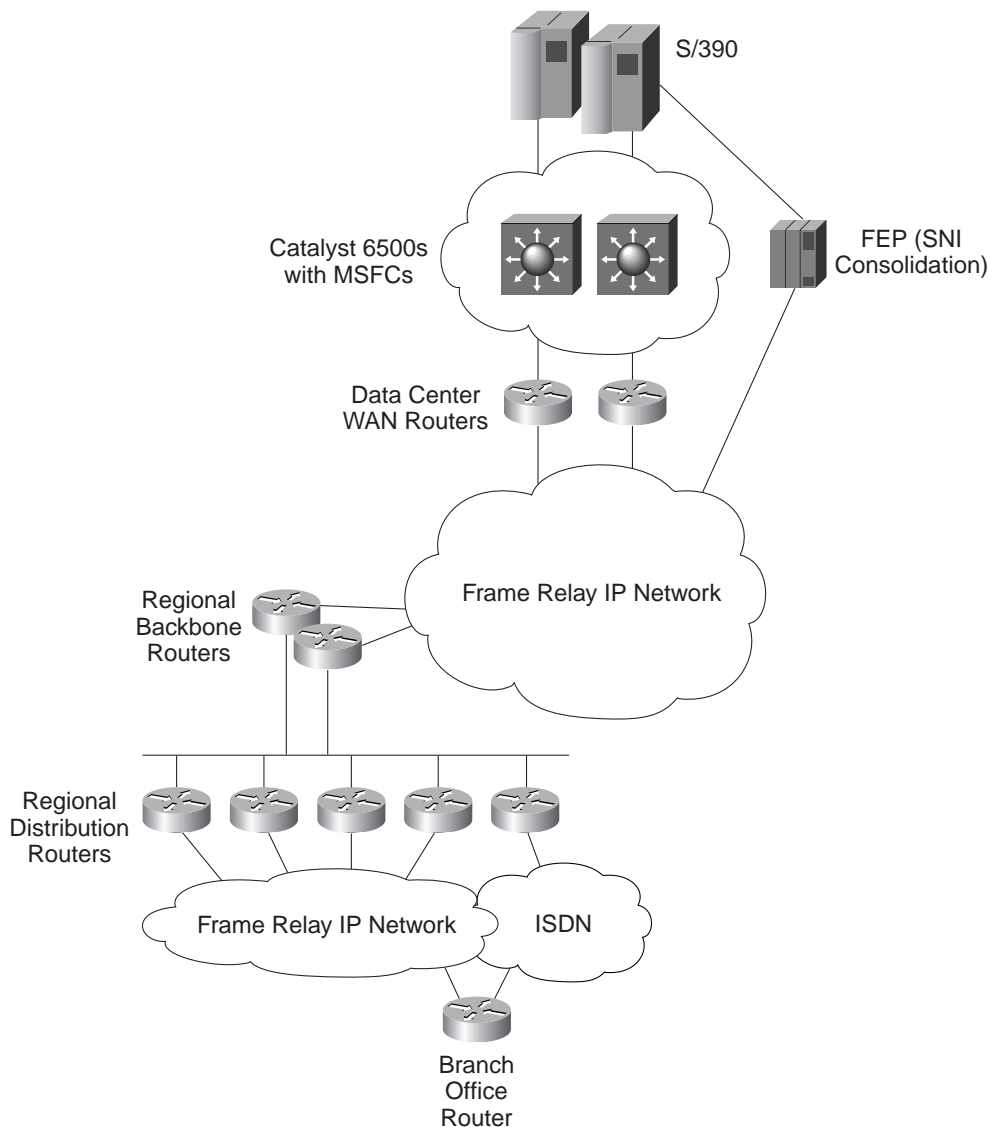
Next the regional distribution routers were installed and configured in the first region. Each distribution layer router would support 50 remote branches with ISDN backup. All of the WAN connections were installed and prepared for later branch network cut-over. This step also involved modifying the old IBM 2216 regional router to establish connectivity to the new network. Each of the 10 regions would be configured much the same way.

Figure 4-4 Convert the Branches and then the Regions



Next, each new Cisco 2600 Series branch router was installed and the branch LAN connection moved over from the old IBM 2210 router. Branch router connectivity to the data center now traveled over the new network while connectivity to branches not yet migrated flowed through the connection to the old IBM 2216 at the regional site.

Figure 4-5 Decommission the Regional IBM 2216 Routers



As each of the regions completed their branch migrations, the IBM 2216s at the regions were decommissioned. The result was a new network using SNASw EE at the branch with improved performance and lower operating cost because of the redesigned Frame Relay network (before, many more PVCs were necessary to support the partial mesh APPN network).

The Old Network versus the New Network

Figures 4-1 through 4-5 depict the network before, during, and after the migration. Figure 4-5 depicts the final newly designed network. The customer achieved all of the goals that were established during design planning sessions. The new IP backbone network provided reliable, efficient data transport. As business needs arise, various QoS and queuing mechanisms can be employed to classify traffic.

The network provided efficient transport of SNA-application traffic using ANR routing at the edge. It also provided RTP connections using Responsive Mode adaptive rate-based flow control (ARB-2) over HPR/IP between SNASw endpoints and between SNASw and the data center hosts. This arrangement resulted in high traffic throughput and system redundancy.

The IP-based data center connectivity resulted in a very high-performance connection to the sysplex environment. Required availability and flexibility was accomplished with redundant switches and routers, IBM OSA-Express adapters, and OSPF routing.

The end result was a network that allowed this large insurance company to better serve its customers and to roll out new applications that would increase productivity and service levels at a rate not possible before with the old network design.

Scenario 2—Data Center and Remote Branch Migration from APPN NN and DLSw to SNASw EE

The organization was a provider of fully integrated local, long distance, and Internet services for international customers in more than 65 countries. The company offered virtual private network (VPN) solutions as well as security, customer care, Web hosting, multicasting, and e-commerce services.

The data center had more than 40 LPARS distributed over four geographically, separate centers. The host traffic consisted of 75 percent SNA/APPN data supported on a dedicated DS-3 network.

Business Requirements

The organization had three business objectives: improve network availability, reduce total network costs, and improve overall network performance. The business requirements that were derived from these objectives were:

- Enable dynamic nondisruptive rerouting of SNA session traffic
- Reduce or eliminate DLSw routers between the remote sites and data center mainframes
- Consolidate equipment and increase network scalability by eliminating APPN NN/DLUR routers at the aggregation locations
- Remove the resource-intensive DLSw function from the ABR routers to enable the ABR routers to accommodate VoIP traffic workload
- Exploit efficient, reliable, direct IP routing to each data center for optimal SNA application access

Network Analysis

To create the new network design, the engineers used a series of questions to help analyze the current network environment. These design questions were briefly discussed in the previous section.

- *Is APPN required?*—As in the previous case study, APPN and APPN/HPR were already enabled in the IBM CS/390 hosts. The major goals of the new network were to improve network availability, enable dynamic nondisruptive rerouting of SNA session traffic over IP, reduce the total costs of networking, and eliminate APPN NN/DLUR and DLSw+ on routers between remote sites and data center hosts. The customer also wanted to improve overall network performance by consolidating the entire network infrastructure over a single IP backbone network (versus dedicated networks for SNA and IP).
- *How do we design for maximum availability and scalability?*—A large part of the customer's design approach was to deploy dynamic routing protocol on the IBM S/390 mainframes. EE between the hosts, with a phased-in deployment to remote branch sites, addressed availability and nondisruptive session path switching when network failures occurred. The customer's deployment of SNASw BX to replace existing APPN NN/DLUR for support of peripheral independent SNA devices addressed the absolute requirement for the new network to be able to scale.

- *How should we connect the Cisco data center routers to the IBM hosts?*—The customer's decision to deploy EE between hosts (EBN) and SNASw out to remote sites for peripheral SNA device support resulted in an IP data center Cisco router solution providing WAN edge and IP dynamic routing transport (SNA transport using DLSw+ or RSRB in data center routers was no longer required).
- *What should we use for transporting SNA traffic from remote sites to the data center?*—EE was chosen as the SNA over IP transport mechanism.
- *Where should we place SNA features such as DLSw+ and SNASw?*—The customer decided to replace remote DLSw+ routers with the SNASw EE feature.
- *How will we manage the integrated network?*—The customer continued to approach network management using existing host NetView platforms.

Design Rationale

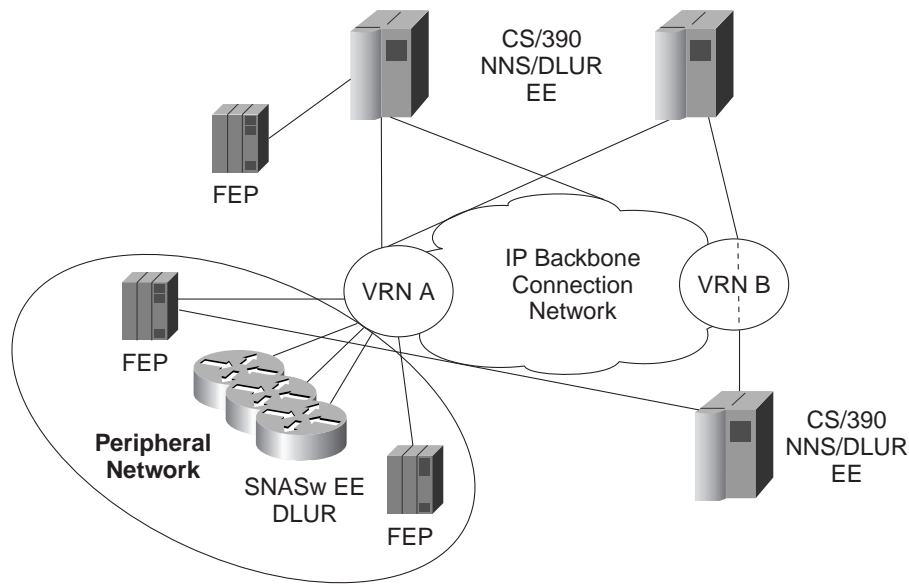
The customer's approach for migrating to EE was initially to implement EBN between mainframes to replace SNI and to IP-enable host mainframe connections. This allowed the customer to become more familiar with using EE in a controlled data center environment and allowed them to address becoming APPN EE-enabled throughout the enterprise. As part of this effort they upgraded the host S/390 operating system to the latest versions of OS/390 and CS/390 and implemented OSPF on their hosts for dynamic routing. One of the customer's biggest requirements was to implement and further extend connection network VRN support to the EE IP network.

The Migration

This migration involved six key steps:

- Step 1. Extend SNASw BX and EE to remote branches to support peripheral DLUR connections
- Step 2. Replace WAN channel-extended and host-to-host DLSw/RSRB
- Step 3. Migrate mainframe-to-mainframe connections to HPR/IP using EE
- Step 4. Implement the EE connection network (VRN)
- Step 5. Map SNA sessions to EE VRN connection network using SNA COS
- Step 6. Implement dynamic routing protocol (OSPF) on the IBM S/390 mainframes

Figure 4-6 Implement the EE Connection Network (VRN)



The Old Network versus the New Network

The new network design, as shown in Figure 4-6, simplified the architecture by creating an IP network infrastructure from the mainframe for transport of both SNA and IP traffic. As a result of this migration, batch transfers achieved a five-fold improvement in performance because of the significantly higher data throughput achieved. In addition, the network had a measurable improvement in network availability as well as simplified network management.

The EE function that was extended to support peripheral network SNA end user traffic over IP resulted in improved session stability. In the future, the customer can take advantage of any new technological changes in the IP WAN network, such as AVVID, without impacting the SNA data transport. The organization also saved \$9.6 million by decommissioning its DS-3 network that was dedicated to handling SNA traffic alone.

Scenario 3—Data Center and Remote Branch Migration from IBM 950 NNs and APPN (PSNA) to SNASw

The organization was a computing provider for savings banks in Germany that had 86 regional branch offices. The core of the network was the computing center, which was operated at four different locations: Duisburg, Cologne/Junkersdorf, Cologne/Gremberghoven, and Mainz. The host SNA applications primarily consisted of CICS, IMS, and proprietary Internet banking applications.

The current network had approximately 3000 Cisco routers (mostly Cisco 4700 Series routers at the regional branch bank offices and Cisco 2504s at customer remote locations). Banking clients of this customer used AS/400s for file transfer, along with service support using OS/2, AIX SNA servers, and SDLC-attached ATM cash machines. The organization had approximately 80-90 regional branch banking offices. Each remote branch bank office had two SNASw routers for redundancy (Cisco 4700 Series routers with 64 MB DRAM memory).

The infrastructure was composed of two SNA network segments, with the SNASw routers split evenly between the two. Each network segment had both primary and backup NN/DLUS servers running on the S/390 host enterprise servers.

Business Requirements

The organization had two business objectives: improve network availability and improve network scalability.

The business requirements that were derived from these objectives were:

- Migrate off installed Cisco APPN NN (PSNA) and IBM NN platforms
- Eliminate APPN broadcast traffic and locate storms from the WAN
- Migrate the backbone network between data centers from SNI to EE EBN
- Migrate the 86 remote branches to EE

Network Analysis

The customer started with a hierarchical SNA network with PU 2.0, 2.1, and LU 6.2 supported by local NCPs (IBM 3745 FEPs) and approximately 50 remote FEPs for boundary function support. Network transport for SNA traffic was supported using Layer 2 RFC 1490 Frame Relay and X.25.

The customer's APPN migration started midyear in 1998 with the deployment of IBM 950 Controller NN servers and 170 NN servers running on Cisco 4700 Series routers with Cisco APPN NN (PSNA) code. Because of network instability and scalability issues encountered as a result of deploying such large numbers of APPN NNs during the initial rollout phase, they decided to migrate to SNASw for BrNN and DLUR support. (The customer also addressed and resolved a number of Y2K issues with some of the older X.25 devices during the migration to SNASw.)

Network management was accomplished through Tivoli NetView for OS/390 plus Internetwork Status Monitor (ISM) on the Cisco 4700 Series routers. Each router was configured with an ISM focal point PU, which allowed them to issue RUNCMD commands via host NetView CLIST control.

Design Rationale

The SNA networks were joined by channel-to-channel connections between DLUS border nodes. The networks had six IBM 950 NN servers channel-attached to their CS/390 hosts. The WAN from the IBM 950s to the SNASw routers was Frame Relay IETF (RFC 1490). Each SNASw router had a single interface to the Frame Relay network with two Frame Relay data-link connection identifiers (DLCIs) connected to each of the IBM 950s. To provide for WAN backup, each SNASw router also had an Integrated Services Digital Network (ISDN) connection to a Cisco 2500 Series router Token Ring-attached to the IBM 950s.

Each SNASw router had three APPN host links. Two of the links were using SRB over Frame Relay (one for each Frame Relay DLCI), and one from a SNASw VDLC port over DLSw+ via ISDN to a TIC connected to the IBM 950 for backup. The ISDN was kept down by configuring no DLSw+ keepalives and by the use of suitable TCP filters. The link via the DLSw+ ISDN backup connection was a low-priority transmission group because it was only used when the other links were down. Links over the ISDN cloud were defined with a lower-bandwidth and higher-cost factor than the Frame Relay links; therefore, no CP-to-CP sessions ran across the ISDN link as long as the Frame Relay connections were active (because the DLSw+ peer for ISDN backup was configured without any keepalive interval).

On the branch banking client user side, each Cisco 4700 Series SNASw router supported two Token Ring interfaces. One Token Ring interface was used by the customer for IP network management. The other Token Ring interface supported SNA client connection using SRB. Each SNASw DLSw+ router was configured as promiscuous so that it could support remote peer DLSw+ connections without having to explicitly configure the DLSw+ remote peers.

Most user traffic was LLC2 either locally bridged or over DLSw+ to SNA servers (most sites had a relatively small number of PUs with many LUs). For the IP network the customer utilized OSPF and Border Gateway Protocol (BGP) as dynamic routing protocols.

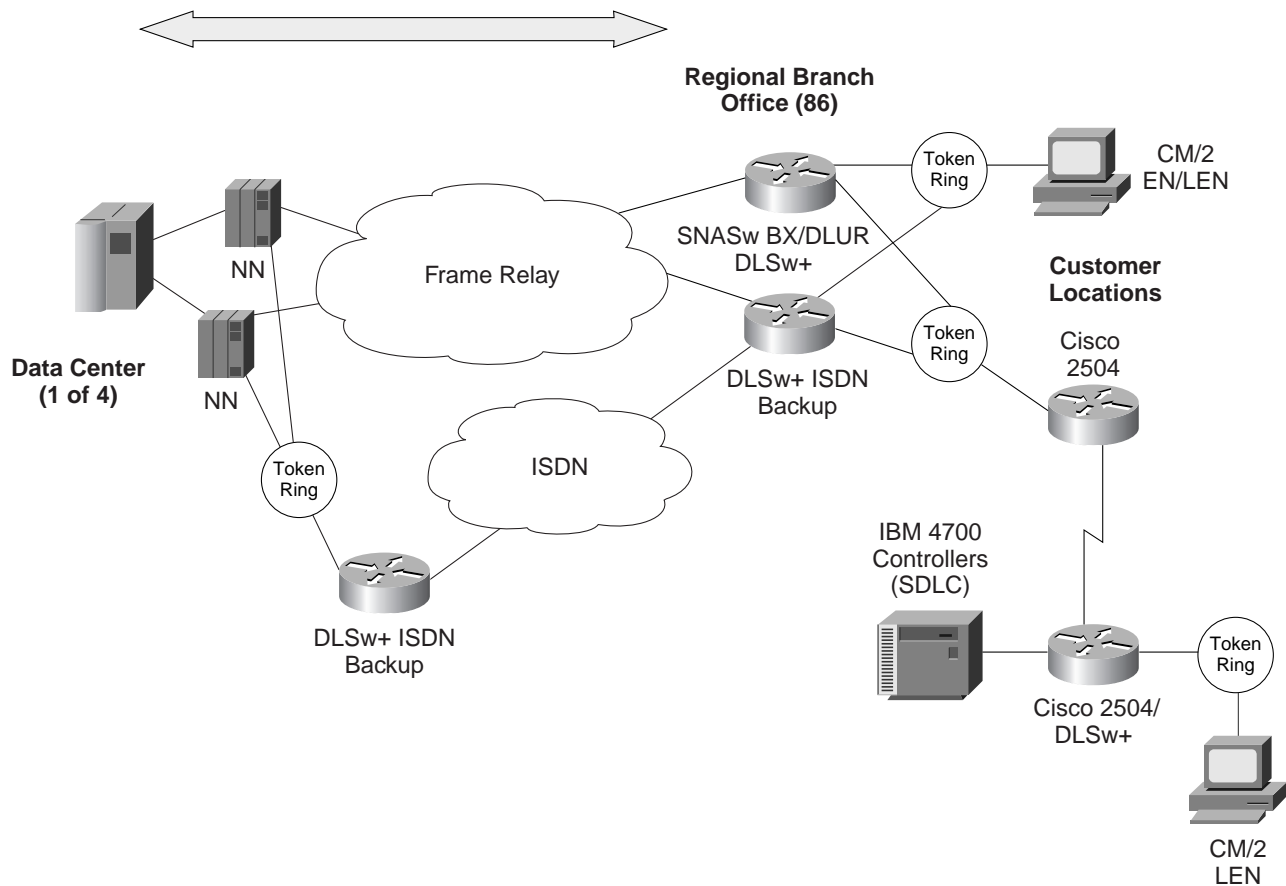
Each SNASw router also had an ISM focal point defined for network management using CiscoWorks Blue (the customer built extensive NetView CLISTS and used RUNCMD commands to operate and monitor their network).

The Migration

The plan involved three phases:

- Phase 1: Migrate regional branch bank offices from APPN NN to SNASw BX/DLUR (see Figure 4-7)
- Phase 2 (future): Migrate data center backbone network from SNI (using IBM FEPs) to HPR/IP EE (EBN)
- Phase 3 (future): Migrate remote regional banks from HPR over LLC2 (Frame Relay RFC 1490 transport) to HPR/IP using SNASw EE

Figure 4-7 Migrating Regional Banks to BX/DLUR





The Old Network versus the New Network

Figure 4-7 depicts the network after the Phase 1 migration. The organization achieved all of the goals that it set during the planning and design phase. The new network addressed their immediate APPN network scalability requirements by eliminating large numbers of APPN NNs at aggregation points at the regional bank branch offices. However, the resulting network after the migration from APPN NN to SNASw Bx was still native APPN/HPR over Layer 2 Frame Relay RFC 1490 transport.

Currently the customer is working with Cisco and IBM to develop plans for migrating to EE later this year. Going forward the customer plans to eliminate the IBM 950s and implement EE EBN support for mainframe-to-mainframe links (the customer is in the process of upgrading the host operating system to OS/390 and CS/390 V2R8). They also plan to upgrade to the latest IBM mainframes that support IBM OSA-Express and to install Cisco Catalyst 6500 Gigabit Ethernet LAN switches in the campus. They are considering upgrading their remote site Cisco 4700 Series routers to Cisco 3640/3660 Series routers to better meet future multiservice (data, voice, video) network IP infrastructure requirements.

Glossary

Advanced Peer-to-Peer Networking—See APPN.

AIW—(APPN Implementers Workshop) The AIW is an industry-wide consortium of networking vendors that develops APPN and SNA-related standards and facilitates high-quality, fully interoperable APPN and SNA internetworking products.

ANR—(Automatic Network Routing) The mode HPR uses to route session traffic between nodes that support RTP functions for HPR. ANR provides point-to-point transport between nodes.

APPN—(Advanced Peer-to-Peer Networking) An extension to SNA that features the following: (1) greater distributed network control that avoids critical hierarchical dependencies, thereby isolating the effects of single points of failure; (2) dynamic exchange of network topology information to foster ease of connection, reconfiguration, and adaptive routing selection; (3) dynamic definition of network resources; and (4) automated resource registration and directory lookup. APPN extends LU 6.2 peer orientation for end-user services to network control and supports multiple LU types including LU 2, LU 3, and LU 6.2.

ARB—(adaptive rate-based) A rate-based congestion and flow control algorithm designed to let APPN HPR RTP connections make efficient use of network resources by providing a congestion avoidance and control mechanism. The basic approach to the algorithm is to regulate the input traffic of an RTP connection based on conditions in the network and conditions at the partner RTP endpoint. When the ARB algorithm detects that the network or the partner endpoint is approaching congestion, ARB reduces the rate at which traffic on an RTP connection is allowed to enter the network until congestion conditions go away.

basic transmission unit—See BTU.

Branch Extender—See BX.

BrNN—(Branch Network Node) See BX.

BTU—(basic transmission unit) In SNA, the unit of data and control information passed between path control components. A BTU can consist of one or more path information units.

channel-attached router—Any Cisco router that is connected to a mainframe via a channel connection using either the CIP or CPA.

BX—(Branch Extender) A function of SNASw that enhances the scalability and reliability of SNA routing nodes by appearing as a NN to downstream EN, LEN node, and PU 2.0 devices while also appearing as an EN to upstream devices. The BX function eliminates APPN topology and APPN broadcast search flows between SNASw nodes and the SNA data hosts in the network.

CBWFQ—(Class-Based Weighted Fair Queuing) An enhanced Cisco QoS functionality that allows different types of network traffic to be prioritized using different map classes.

CICS—(Customer Information Control System) An IBM application subsystem allowing transactions entered at remote terminals to be processed concurrently by user applications.

CIP—(Channel Interface Processor) A Cisco interface processor for the Cisco 7000 and 7500 Series routers that provides ESCON or bus and tag channel attachment to the mainframe.

Class of Service—See COS.

CLSI—(Cisco link services interface) An interface architecture that allows various Cisco data-link user Cisco IOS features such as DLSw+ and VDLC to interface with and access the services of data-link control protocol stacks (such as SDLC).

CMCC—(Cisco Mainframe Channel Connection) Any of the Cisco router CIP and CPA feature cards (interface processors or port adapters) that allow a user to establish a channel connection between the router and a mainframe.

CMCP+—(Cisco MultiPath Channel Plus) CMPC+ enables High Performance Data Transfer (HPDT). It allows TCP/IP connections to the host through CMCC adapters, using either the TCP/IP stack or the High Speed Access Services (HSAS) IP stack.

CN—(connection network) A representation within an APPN network of shared-access transport facilities (SATFs), such as Token Ring, that allows nodes identifying their connectivity to the SATF by a common virtual routing node (VRN) to communicate without having individually defined connections to one another.

connection network—See CN.

CNN—(composite network node) A node representing a group of nodes that appear as one APPN or LEN node to other nodes in an APPN network. For example, a subarea network consisting of a VTAM host and some NCPs is a multiple-node network, but when connected to an APPN node, it appears as *one* logical APPN or LEN node.

COS—(Class of Service) A set of characteristics such as specific transmission priority, level of route reliability, and security level used to construct a route between session partners. The COS is derived from a mode name specified by the initiator of a session.

CP—(control point) In SNA networks, an element that identifies the APPN networking components of a PU 2.1 node, manages device resources, and provides services to other devices. In APPN, CPs are able to communicate with logically adjacent CPs using CP-to-CP sessions.

CPA—(Channel Port Adapter) A Cisco port adapter for the Cisco 7200 Series routers that provides ESCON or bus and tag channel attachment to the mainframe.

CSNA—(Cisco SNA) An application that provides support for SNA protocols to the IBM mainframe from Cisco 7500 Series CIP2 and Cisco 7200 Series CPA platforms.

DLCI—(data-link connection identifier) A value that specifies a PVC or SVC in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant (connected devices might use different values to specify the same connection). In the LMI extended specification, DLCIs are globally significant (DLCIs specify individual end devices).

DLSw—(data-link switching) An interoperability standard, described in RFC 1795 and 2166, that provides a method for forwarding SNA and NetBIOS traffic over TCP/IP networks using data-link layer switching and encapsulation. DLSw uses SSP instead of SRB, eliminating the major limitations of SRB, including hop-count limits, broadcast and unnecessary traffic, time outs, lack of flow control, and lack of prioritization schemes.

DLSw+—(Data-Link Switching Plus) Cisco's implementation of DLSw, which includes significant scalability, availability, transport flexibility, and load-balancing enhancements over the DLSw version 1 (RFC 1795) and version 2 (RFC 2166) standards.

DLU—(dependent LU) An LU that requires assistance from an SSCP in IBM CS/390 to initiate an LU-to-LU session.

DLUR—(Dependent LU Requester) A feature of APPN that allows traditional dependent SNA subarea traffic to be routed over the APPN network.

DLUS—(Dependent LU Server) The server half of the DLUR/DLUS enhancement to APPN. The DLUS component provides SSCP services to DLUR nodes over an APPN network.

DSPU—(Downstream Physical Unit) A software feature that enables the router to function as a PU concentrator for SNA PU 2.0 nodes.

EBN—(extended border node) An APPN node type that allows the connection of NNs with different NETIDs and allows session establishment between LUs in different NETID subnetworks that need not be adjacent.

EE—(Enterprise Extender) A function of SNASw that offers SNA High Performance Routing (HPR) support directly over IP networks, utilizing connectionless UDP transport.

EN—(end node) An APPN end system that implements the PU 2.1, provides end-user services, and supports sessions between local and remote CPs. ENs are not capable of routing traffic and rely on an adjacent NN for APPN services.

Enterprise Extender—see EE.

ESCON—(Enterprise Systems Connection) A data processing environment having a channel-to-control unit input/output interface using optical cables as the transmission media.

FDDI—(Fiber Distributed Data Interface) A LAN standard, defined by ANSI X3T9.5, specifying a 100-Mbps token-passing network using fiber-optic cable, with transmission distances of up to 2 km. FDDI uses a dual-ring architecture to provide redundancy.

FEP—(front-end processor) A device or board that provides network interface capabilities for a networked device. In SNA, an FEP is typically an IBM 3745 device.

FRAS—(Frame Relay Access Support) A Cisco IOS Software feature that allows branch SNA devices to connect directly to a central site FEP over a Frame Relay network.

FST—(Fast Sequenced Transport) A high-performance DLSw+ encapsulation option used over higher-speed links (256 kbps or higher) when high throughput is required.

HDLC—(High-Level Data Link Control) A bit-oriented synchronous data link layer protocol developed by ISO. Derived from SDLC, HDLC specifies a data encapsulation method on synchronous serial links using frame characters and checksums.

High Performance Routing—See HPR.

HPR—(High Performance Routing) An addition to APPN that enhances data-routing performance and session reliability.

ICN—(interchange node) A standalone APPN CS/390 node or a CNN. The ICN routes sessions from APPN nodes into and through the subarea network using subarea routing, without exposing the subarea implementation to the APPN part of the network.

IETF—(Internet Engineering Task Force) A task force consisting of more than 80 working groups responsible for developing Internet standards.

Intermediate Session Routing—See ISR.

IPM—(Internetwork Performance Monitor) A Cisco workstation-based network management product that provides data about response times between devices.

ISDN—(Integrated Services Digital Network) A communication protocol, offered by telephone companies, that permits telephone networks to carry data, voice, and other source traffic.

ISM—(Internetwork Status Monitor) A Cisco mainframe-based network management product that allows users to manage their Cisco routers from their mainframe network management application (Tivoli NetView for OS/390). ISM enables NetView operators to have full visibility of a Cisco router network from a single NetView console regardless of whether that router network is routing SNA traffic.

ISR—(Intermediate Session Routing) A type of routing function within an APPN NN that provides session-level flow control and outage reporting for all sessions that pass through the node but whose endpoints are elsewhere.

LEN node—(low-entry networking node) A capability of nodes to attach directly to one another using basic peer-to-peer protocols without APPN to support parallel LU 6.2 sessions between LUs.

LLC2—(Logical Link Control, type 2) A connection-oriented SNA LLC-sublayer protocol.

low-entry networking node—See LEN node.

LPAR—(logical partition) A physical IBM S/390 or zSeries mainframe divided into multiple logical partitions.

LU—(logical unit) A type of addressable unit in an SNA network. The LU is the port through which the end user accesses both the SSCP-provided services and communicates with other LUs at other nodes.

MAC—(Media Access Control) The lower of the two sublayers of the data link layer defined by the IEEE. The MAC sublayer handles access to shared media, such as whether token passing or contention will be used.

maximum transfer unit—See MTU.

MIB—(Management Information Base) A database of network management information that is used and maintained by a network management protocol such as SNMP.

MSFC—(Multilayer Switch Feature Card) A card that provides Cisco IOS Software Layer 3 services on Cisco Catalyst 6500 Series switches.

MTU—(maximum transfer unit) The maximum packet size in bytes that a particular interface can support.

network node server—See NN server.

NN—(network node) An APPN node type that provides full distributed directory and routing services for all LUs that it controls. These LUs can be located on the APPN NN itself or on one of the adjacent LEN nodes or APPN ENs for which the APPN NN provides NN services. Jointly with the other active APPN NNs, an APPN NN is able to locate all destination LUs known in the network.

NN server—(network node server) An APPN NN that provides network services for its local LUs and client ENs.

OSA-Express—(Open Systems Adapter-Express) An IBM mainframe NIC used to attach a Cisco Catalyst 6500 Series Gigabit Ethernet switch to an IBM S/390 or zSeries host.

OSPF—(Open Shortest Path First) A link-state, hierarchical routing algorithm that features least-cost routing, multipath routing, and load balancing.

PBN—(peripheral border node) An APPN node type that enables the connection of NNs with different NETIDs and allows session establishment between LUs in different, adjacent subnetworks.

PPP—(Point-to-Point Protocol) A protocol that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits.

PSNA—(Portable Systems Network Architecture) The first-generation Cisco APPN NN platform in the Cisco IOS Software.

PU—(physical unit) A type of addressable unit in an SNA network. Each node in the network has a PU, which provides services to control the physical configuration and the communication system resources associated with the node and to collect maintenance and operational statistics.

QLLC—(Qualified Logical Link Control) A data link layer protocol defined by IBM that allows SNA data to be transported across X.25 networks.

QoS—(quality of service) The measure of performance for a transmission system that reflects its transmission quality and service availability.

quality of service—See QoS.

Responsive Mode ARB—The second-generation enhanced HPR flow control algorithm used for SNA transport over HPR/IP (EE) networks.

RFC—(Request for Comments) A document series used as the primary means for communicating information about the Internet. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications such as Telnet and FTP, but some are humorous or historical. RFCs are available online from numerous sources.

RIF—(routing information field) A field in the IEEE 802.5 header that is used by a source-route bridge to determine the Token Ring network segments through which a packet must transit. A RIF is made up of ring and bridge numbers as well as other information.

RSRB—(remote source-route bridging) Cisco's first technique for connecting Token Ring networks over *non-Token Ring* WAN network segments.

RTP—(Rapid Transport Protocol) A connection-oriented, full-duplex protocol designed to transport data in high-speed networks. HPR uses RTP connections to transport LU-to-LU and CP-to-CP session traffic. RTP provides reliability, in-order delivery, segmentation and reassembly, and adaptive rate-based flow/congestion control. Because RTP provides these functions on an end-to-end basis, it eliminates the need for these functions on the link level along the path of a connection.

SATF—(shared access transport facility) A shared-access medium that allows for dynamic direct connectivity between any pair of link stations attaching to the facility.

SDLC—(Synchronous Data Link Control) An SNA data link layer bit-oriented, full-duplex serial communications protocol.

SNA—(Systems Network Architecture) The IBM architecture that defines the logical structure, formats, protocols, and operational sequences for transmitting information units through, and controlling the configuration and operation of, networks. The layered structure of SNA allows the ultimate origins and destinations of information (the users) to be independent of and unaffected by the specific SNA network services and facilities that are used for information exchange.

SNASw—(SNA Switching Services) A feature within the Cisco IOS Software that provides SNA routing or "session switching" for PU 2.0, PU 2.1 (LEN node), and APPN EN devices over ISR or HPR data-link controls.

SNI—(SNA network interconnection) The connection by gateways of two or more independent SNA networks to allow communication between SNA LUs in those networks.

SNMP—(Simple Network Management Protocol) A network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices and to manage configurations, statistics collection, performance, and security.

SRB—(source-route bridging) A method of bridging originated by IBM and popular in Token Ring networks. In an SRB network, the entire route to a destination is predetermined, in real time, before data is sent to the destination.

SSCP—(System Services Control Point) The SNA architectural component within a subarea network for managing the configuration, coordinating network operator and problem determination requests, and providing directory services and other session services for end users of a network. Multiple SSCPs, cooperating as peers with one another, can divide the network into domains of control, with each SSCP having a hierarchical control relationship to the PUs and LUs within its own domain.

sysplex—A set of MVS or OS/390 systems communicating and cooperating with each other through certain multisystem hardware components and software services to process customer workloads. This term is derived from "system complex."

System Services Control Point—See SSCP.

Systems Network Architecture—See SNA.

TAC—(Technical Assistance Center) Cisco's world-class technical support center.

TDU—(topology database update) A message about a new or changed link or node that is broadcast among APPN NNs to maintain the network topology database. TDUs are fully replicated in each NN in an APPN network.

TIC—(Token Ring interface coupler) An adapter that can connect an IBM FEP to a Token Ring network.

UDP—(User Datagram Protocol) A connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.

VDLC—(Virtual Data Link Control) A function within the Cisco IOS Software that provides communication between two software components that both use Cisco link services (CLS).

Virtual Data Link Control—See VDLC.

VoIP—(Voice over IP) A technology that enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network.

VRN—(virtual routing node) A representation of a node's connectivity to an APPN connection network defined on an SATE.

WFQ—(Weighted Fair Queuing) A congestion management algorithm that identifies conversations (in the form of traffic streams), separates packets that belong to each conversation, and ensures that capacity is shared fairly between these individual conversations. WFQ is an automatic way of stabilizing network behavior during congestion and results in increased performance and reduced retransmission.

XCA—(external communication adapter) A communication adapter used by the CIP and CPA to allow one host subchannel to support thousands of SNA PUs to a CS/390 host. An XCA also can define the IP port (the connection to the adjoining CS/390 host TCP/IP stack) that CS/390 will use for EE connections.

APPN Components and Features

APPN Technology Overview

An SNA node is a set of hardware and associated software components that implement the functions of the seven SNA architectural layers. Although all seven layers are implemented within a given node, nodes can differ based on their architectural components and the sets of functional capabilities they implement.

Every node defined by SNA contains a CP. A PU 5 subarea node contains SSCP. An SSCP activates, controls, and deactivates SNA resources in the network. The CP in a PU 2.1 node is called simply a CP. In general, a CP manages the network resources such as PUs, LUs, and sessions.

Advanced Communications Function (ACF)/VTAM in CS/390 contains a CP that manages SNA devices in its span of control (that is, within its domain). In general, all devices within the VTAM domain must be configured to that VTAM, but VTAM can dynamically find resources “owned” by another VTAM. VTAM is responsible for establishing all sessions and for activating and deactivating resources in its domain. In this environment, resources are explicitly predefined, thereby eliminating the requirement for broadcast traffic and minimizing header overhead.

A PU provides services needed to manage a specific type of device and any resources associated with the device. A PU is composed of software, hardware, and microcode. It represents a set of functions. If software such as IBM's VTAM performs these functions, then the software represents the PU. A PU can also be a piece of hardware, such as a FEP. For some devices (for example, an IBM 3174 controller), the PU function is included within the terminal's microcode.

The following list includes the PU types within the SNA architecture:

- PU 5 (mainframe)
- PU 4 (FEP)
- PU 2.0
- PU 2.1 (LEN)
- PU 1

APPN Components and Node Types

The components and node types that make up the APPN architecture are as follows:

- NN
- NN server
- EN
- LEN node

- Composite network node (CNN)
- CP
- BrNN
- HPR node
- ICN
- DLUR/DLUS
- Dependent and independent LUs
- Border node
- Migration data host
- Transmission group
- VRN
- Central directory server (CDS)

NN

NNs make up the backbone routers in an APPN network. Together with the links that interconnect them, the NNs form the intermediate routing network of an APPN network. The NNs connect the ENs to the network and provide resource location and route selection services for them.

A NN provides the following functions for ENs and LEN nodes:

- Distributed directory services
- Topology database exchanges with other APPN NNs
- Session services for local LUs and client ENs
- Intermediate routing services

The original APPN architecture was defined so that NNs maintain both local and network topology databases. When an EN requests a session setup for a pair of resources, the NN will first look in its directory to see if it knows the location of the destination. If it does, session setup can proceed. If it does not know the location of the destination, the NN will send a broadcast throughout the network to locate the destination. When the destination is found, the NN adds information about the destination to its directory, selects a session path to meet the COS defined in the session setup request, and then tells the EN to complete session setup.

NN Server

A NN server is a NN that provides resource location and route selection services to the LUs that it serves. These LUs can be in the NN itself or in the client ENs. A NN server uses CP-to-CP sessions to provide network information for session setup in order to support the LUs on served APPN ENs. In addition, LEN ENs can take advantage of the services of the NN server.

EN

An EN is located on the periphery of an APPN network. An EN obtains full access to the APPN network through one of the NNs to which it is directly attached and serves as its NN server. There are two types of ENs: APPN ENs and LEN ENs. The APPN ENs provide limited directory and routing services only for local resources and support APPN protocols through a NN server. A LEN EN is a LEN that is connected to an APPN NN. Although the LEN EN lacks the APPN extensions, it is provided services through its NN.



LEN Node

A LEN node provides peer-to-peer connectivity to other LEN nodes, APPN ENs, or APPN NNs. A LEN node requires that all network accessible resources, either controlled by the LEN node itself or on other nodes, be defined at the LEN node.

Unlike APPN ENs, the LEN node cannot establish CP-to-CP sessions with an APPN NN. A LEN node therefore cannot register resources at a NN server. Nor can it request a NN server to search for a resource or to calculate the route between itself and the node containing a destination resource. It does, however, use the distributed directory and routing services of an adjacent NN indirectly. It does this by predefining remote LUs, owned by nonadjacent nodes, with the CP name of an adjacent APPN NN.

CNN

A CNN is a PU 5 node (VTAM) along with its subordinate PU 4 nodes (NCP, which runs in the FEP). Unlike VTAM, which can be a standalone NN or EN, NCP cannot provide this capability. Therefore, working together, they can represent a single NN.

CP

The CP in a PU 2.1 or PU 5 node is simply called a CP. Like other CPs, it has a number of functions: it can activate locally attached links, interact with a local operator, and manage local resources. It can also provide network services, such as partner location and route selection, for local LUs in a subarea network.

A CP for a LEN node does not communicate with a CP in another node. It communicates only with other components in its own node for controlling local resources and for aiding local LUs in establishing LU-to-LU sessions.

An APPN EN CP participates in CP-to-CP sessions with the CP in an adjacent NN server. Two parallel sessions using LU 6.2 protocols are established between the partner CPs. The APPN EN does not establish CP-to-CP sessions, however, with any adjacent LEN node or APPN ENs. If it is attached to multiple APPN NNs, the APPN EN chooses only one of them to be its active server; it does not establish CP-to-CP sessions with more than one NN at a time. (It can, however, route LU-to-LU sessions through any adjacent NN as determined by route selection criteria.)

BrNN

The SNASw BrNN feature eliminates the need for the full NN functionality in the network and provides a more scalable solution by effectively eliminating SNA topology and broadcast search traffic from the network.

HPR Node

An HPR node is an APPN node that has implemented the optional HPR functions. An HPR node can be an APPN EN or an APPN NN. HPR is an enhancement to APPN that provides improved network performance and reliability. HPR replaces ISR with two elements: ANR and RTP.

APPN HPR is capable of SNA session recovery. In case of a route failure, the RTP layer selects an alternate path and reroutes SNA packets. Before HPR, neither SNA nor APPN was capable of LU-to-LU session recovery. Within HPR, RTP provides end-to-end processing, reordering, and delivery of frames, which produces nondisruptive route switching and flow control. On the other hand, ANR provides node-to-node connectionless source routing. These features of HPR create a protocol that is similar in functionality to TCP/IP.

ICN

The most common migration step for a multidomain network is to move the CMC VTAM subarea host to an ICN. An ICN combines the function of a subarea node and a NN. It implements APPN NN, SSCP, and cross-domain resource manager (CDRM) functionality. An ICN is located on the border of an APPN network and a SNA subarea network. The node converts session requests between subarea SNA and APPN protocols, thus enabling sessions to be established between applications residing on an APPN VTAM host to LUs residing in the subarea network.

An ICN can own and activate NCPs. It communicates network control data by using SSCP-to-SSCP sessions with other subarea nodes and CP-to-CP sessions with other APPN nodes. To enable it to participate in the subarea network, it is defined with a unique subarea number and requires subarea path definition statements. An ICN can be connected to other APPN nodes, LEN nodes, and subarea nodes.

DLUR/ DLUS

The functionality of the DLUR/DLUS was a direct result of the thousands of SNA devices that still require the master/slave relationship for network connectivity and data delivery over APPN. The DLUR functionality is implemented in an APPN EN or NN (remote to VTAM), whereas DLUS is implemented in a VTAM APPN NN (version 4.2 or higher).

Dependent and Independent LUs

In SNA subarea networks, LUs that reside in peripheral nodes can be dependent (SSCP-dependent) or independent (SSCP-independent). An independent LU is dependent or independent based on the protocols it uses to initiate LU-to-LU sessions.

An independent LU sends a session-activation request to another LU. No SSCP intervention is required. LUs reside in PU 2.1 nodes.

In APPN networks, independent LUs establish sessions directly with partners by sending session-activation requests directly to the partner. Allowing independent LUs to establish direct sessions with other independent LUs (without the intervention of VTAM) is referred to as peer-to-peer communication.

APPN independent LUs can dynamically register themselves to their NN server. ENs issue a Locate request to their NNs to dynamically locate APPN resources. When an EN sends a directory search to a NN, the NN forwards a Locate to all of its NN neighbors, and the neighbors propagate the Locate throughout the network. When the resource is found, the directory is updated, and the EN is informed of the resource location.

When a NN initializes, it exchanges its topology with its directly attached neighbors. Every time a change occurs in the topology database of a NN, the NN propagates the change to all directly attached NNs. In this manner, all NNs have an exact copy of the topology database. From this topology database, each NN builds one or more trees, with itself as the root (there is one tree per COS).

When a session is established, the best path from source to destination is determined. The same path is used for the duration of the session.

A dependent LU uses SSCP-to-LU sessions to send session-initiation requests to the controlling SSCP. The dependent LU depends on the SSCP to conduct the session initiation with the target LU and requires an SSCP-to-LU session. APPN provides support for SSCP-dependent LUs using DLUR/DLUS.

The DLUS is a product feature of a PU 5 (VTAM) NN. The DLUS function enables VTAM to have SSCP services for dependent LUs in remote APPN ENs or NNs; the ENs and NNs act as the DLUR.



Border Node

Base APPN architecture does not allow two adjacent APPN NNs to connect and establish CP-to-CP sessions when they do not have the same network identifier (NETID). The border node is an optional feature of an APPN NN that overcomes this restriction. A border node can connect to an APPN NN with a different NETID, establish CP-to-CP sessions with it, and then allow session establishment between LUs in different subnetworks. Topology information is not passed between the subnetworks. In addition, a border node can also connect to another border node.

Migration Data Host

A migration data host is a VTAM data host that communicates to APPN resources as an APPN EN and communicates to directly attached VTAMs and NCPs using subarea flows. Unlike an ICN, a migration data host does not translate between APPN and subarea protocols, and it cannot own NCPs or provide ISR. A migration data host is often implemented during migration from subarea SNA to APPN because the migration allows data hosts to be accessed from an APPN node or subarea node concurrently.

Transmission Group

A transmission group is the same in APPN terminology as it is in legacy SNA. A transmission group is the set of lines connecting two nodes. The difference between a transmission group in SNA and in APPN is that the current APPN architecture limits a transmission group to a single link, although multilink transmission groups are expected to be implemented in the future. The topology database contains NNs and transmission groups that connect NNs.

VRN

A VRN represents a connection to the SATF, such as Token Ring and FDDI. The SATF and the set of nodes having defined connections to a common VRN are said to comprise a connection network.

CDS

A central directory server is a NN that builds and maintains a directory of resources within the network. The purpose of a CDS is to reduce the number of network broadcast searches to a maximum of one per resource. NNs and ENs can register their resources with a CDS, which acts as a focal point for resource location information.

A CDS can be involved in APPN EN resource registration. An APPN EN registers its resources to improve network search performance. When an APPN EN registers its resources with its NN server, it can request that its NN server also register them with a CDS. Entries in a directory database can be registered, defined, or dynamic.


When a NN receives a search request for a resource that has no location information, the NN first sends a directed search request to a CDS, if there is one. The CDS searches in its directory for information about the location of the resource. If it does not find the location of the resource, the CDS searches ENs in its domain, other NNs and ENs, and, if necessary, the entire network (via a broadcast search). If the resource is still not found, the CDS notifies the NN that originally requested the search that the search is unsuccessful. A central directory client is a NN that forwards directory searches to a CDS.

Enterprise Extender (HPR/IP) Sample Configuration

Overview

This sample configuration illustrates HPR using IP. Figure C-1 shows how SNASw can be used to connect to downstream devices. In this case, a downstream physical unit (DSPU) router is simulating a downstream PU connected to CS/390 through the DLUR. The DSPU router is configured as a downstream PU. The PU connects to the virtual Token Ring interface on the SNASw router by means of SRB over the physical Token Ring interface. The upstream connection to the host is done through IP. The CIP router is configured to run Cisco MultiPath Channel Plus (CMPC+).

Figure C-1 Network Topology



To implement this configuration, you need to have the following host definitions:

- An external communications adapter (XCA) major node in CS/390 with “MEDIUM=HPRIP” defined
- A switch major node for the SNASw CP
- A switch major node for the downstream PU
- A Transport Resource List (TRL)
- A Profile TCP/IP with device name matching CS/390 TRL entry

Note: The downstream devices are not limited to LAN-based connections; therefore, you can use SDLC and so on. You can also use DLSw+ for downstream devices and connect into the SNASw using a VDLC port.

SNASw Configurations

SNASw Router

```
Current configuration:
!
version 12.0
hostname SNASW
!
boot system flash slot0:rsp-a3jsv-mz.120-5.XN
enable password lab
!
ip subnet-zero
!
source-bridge ring-group 100
!
interface Ethernet0/0/0
 ip address 172.18.49.37 255.255.255.128
 no ip directed-broadcast
 no ip route-cache distributed
!
interface TokenRing2/0/2
 no ip address
 no ip directed-broadcast
 no ip route-cache distributed
 ring-speed 16
 source-bridge 200 1 100
 source-bridge spanning
interface Virtual-TokenRing2
description this interface is used to connect in the downstream PU
 mac-address 4000.eeee.0000
 no ip address
 no ip directed-broadcast
 ring-speed 16
 source-bridge 222 1 100
 source-bridge spanning
snasw cpname NETA hostname
snasw port HPRIP hpr-ip Ethernet0/0/0 vnname NETMD.EEJEB
snasw port VTOK2 Virtual-TokenRing2 vnname NETMD.EEJEB
snasw link HPRMVSD port HPRIP ip-dest 172.18.1.41
router eigrp 109
 network 172.18.0.0
 no auto-summary
!
ip classless
line con 0
 exec-timeout 0 0
 transport input none
line aux 0
line vty 0 4
 login
!
end
SNASW#
```

DSPU Router

```
!  
hostname DSPU  
!  
boot system flash  
enable password lab  
!  
ip subnet-zero  
!  
source-bridge ring-group 300  
  
dspu host TOKEN xid-snd 02201002 rmac 4000.eeee.0000 rsap 4 lsap 12  
dspu pool pool_lu host TOKEN lu 2 2  
!  
  
interface TokenRing0/0  
no ip address  
no ip directed-broadcast  
ring-speed 16  
source-bridge 200 1 300  
dspu enable-host lsap 12  
dspu start TOKEN  
!  
line con 0  
exec-timeout 0 0  
transport input none  
line aux 0  
line vty 0 4  
password lab  
login  
!  
end  
  
DSPU#
```

CIP Router

```
Current configuration:
!
version 12.0
hostname CIPRouter
!
enable password lab
!
microcode CIP flash slot0:cip27-6
microcode reload
ip subnet-zero
source-bridge ring-group 80
interface Ethernet0/0
  ip address 172.18.49.17 255.255.255.128
  no ip directed-broadcast
  no ip mroute-cache
interface Channell1/0
  no ip address
  no ip directed-broadcast
  no keepalive
!
interface Channell1/1
  no ip address
  no ip directed-broadcast
  no keepalive
  cmpc E160 92 EETGJEB READ
  cmpc E160 93 EETGJEB WRITE
!
interface Channell1/2
  ip address 172.18.1.42 255.255.255.248
  no ip directed-broadcast
  no ip mroute-cache
  no keepalive
  lan TokenRing 0
    source-bridge 70 1 80
    adapter 0 4000.dddd.aaaa
  tg EETGJEB ip 172.18.1.43 172.18.1.42
router eigrp 109
  network 172.18.0.0
  no auto-summary
!
ip classless
ip route 172.18.1.41 255.255.255.255 172.18.1.43
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0
  exec-timeout 0 0
  password lab
  login
  length 75
  width 114
line vty 1 4
  exec-timeout 0 0
  password lab
  login
!
end
CIPRouter#
```


Host Definitions

```
CISCO.NETMD.VTAMLST(XCAEEJEB)
```

```
-----
EEXCAJ VBUILD TYPE=XCA
EETGJ PORT MEDIUM=HPRIP, X
VNNAME=EEJEB, X
VNGROUP=EEGRPJ, X
LIVTIME=15, X
SRQTIME=15, X
SRQRETRY=9, X
SAPADDR=04
*
EEGRPJ GROUP ANSWER=ON, X
AUTOGEN=(64,L,P), X
CALL=INOUT, X
DIAL=YES, X
DYNPU=YES, X
DYNPUPFX=$E, X
ISTATUS=ACTIVE
```

```
CISCO.NETMD.VTAMLST(EETGJEB)
```

```
-----
EETGJEBV VBUILD TYPE=TRL
EETGJEB TRLE LNCTL=MPC,MAXBFRU=16, X
READ=(4F92), X
WRITE=(4F93)
```

```
PROFILE.TCPIP
```

```
DEVICE IUTSAMEH MPCPTP AUTORESTART
LINK samehlnk MPCPTP IUTSAMEH
;
DEVICE EETGJEB MPCPTP
LINK EELINK2 MPCPTP EETGJEB
;
DEVICE VIPADEV2 VIRT 0
LINK VIPALNK2 VIRT 0 VIPADEV2
;
HOME
172.18.1.43 EELINK2 ; This corresponds to the host-ip-addr for the CIPRouter tg
command.
172.18.1.41 VIPALNK2 ; This corresponds to the ip-dest specified in the SNASW router
link command.
GATEWAY
172.18 = EELINK2 4468 0.0.255.248 0.0.1.40
172.18 172.18.1.42 EELINK2 4468 0.0.255.0 0.0.49.0
;
START IUTSAMEH
START EETGJEB
```

```
VIEW CISCO.NETMD.VTAMLST(SNASWCP) - 01.02 Columns 00001 00072
***** ***** Top of Data *****
==MSG> -Warning- The UNDO command is not available until you change
```

```

==MSG>          your edit profile using the command RECOVERY ON.
000001 *          SNASWITCH CONTROL POINT
000002          VBUILD TYPE=SWNET
000003 *
000004 R7507PU  PU      ADDR=01,ANS=CONTINUE,DISCNT=NO,                      X
000005          PUTYPE=2,ISTATUS=ACTIVE,                                      X
000006          NETID=NETA,CPCP=YES,CONNTYPE=APPN,CPNAME=SNASW,HPR=YES
000007
***** ***** Bottom of Data *****

```

```

VIEW          CISCO.NETMD.VTAMLST(SNASWPUS) - 01.02          Columns 00001 00072
***** ***** Top of Data *****
==MSG> -Warning- The UNDO command is not available until you change
==MSG>          your edit profile using the command RECOVERY ON.
000001 *          SNASWITCH DOWNSTREAM PU
000002          VBUILD TYPE=SWNET
000003 *
000004 DSPU02  PU      ADDR=01,ANS=CONTINUE,DISCNT=NO,                      X
000005          PUTYPE=2,ISTATUS=ACTIVE,                                      X
000006          DLOGMOD=D4C32782,MODETAB=ISTINCLM,USSTAB=USSTCPMF,          X
000007          IDBLK=022,IDNUM=01002                                          X
000008 DSPU02LU LU      LOCADDR=02
***** ***** Bottom of Data *****

```

SNASw Verification Commands

Table C-1 lists some of the commands for verifying that SNASw is running in the router. For a full list, see the *Cisco IOS Bridging and IBM Networking Command Reference*.

Table C-1 SNASw Show and Trace Commands

Command	Description
show snasw session	To see the session and partner details
show snasw dlus	To verify if the DLUS is active
show snasw pu	To verify if the PU is active
show snasw link	To verify which port the link uses, the node type, and if the link is active
snasw dlctrace	To trace frames arriving and leaving SNASw
show dlctrace	To display a trace on the console; alternatively, you can dump the trace onto a server

SNASw Hardware and Software Requirements

Supported Hardware

SNASw is supported on the following hardware platforms:

- Cisco 2500 Series routers
- Cisco 2600 Series routers
- Cisco 3600 Series routers
- Cisco 4000 Series routers
- Cisco 7200 Series routers
- Cisco 7500 Series routers
- Catalyst 5000 Series switches with Route Switch Modules (RSMs)

Supported Software

SNASw supports the following software:

- Cisco IOS Release 12.1 or higher
- IP Plus/SNASw Plus Software (Cisco 2500 and 4500M platforms only)
- Enterprise/SNASw Plus IOS software suites (all platforms)

SNASw EE requires the following host software:

- IBM CS/390 V2R6 with APAR OW36113 or higher (CS/390 V2R7 or higher is recommended because of host IP stack enhancements)

APARs

Software requirements for the host that has BX includes requirements for IBM APARs. SNASw BX works with all IBM supported CS/390 releases and the following IBM APARs:

- OW37548
- OW37549
- OW39559

The additional APARs required for APPN HPR are as follows:

- OW43416
- OW43223
- OW43413

- OW41980
- OW38056
- OW43484

Using Responsive Mode ARB on OS390 V2R7 requires the following APARs:

- OW36968/UW55496
- OW36694/UW56215
- OW38588/UW59095

Supported MIBs, RFCs, and Standards

SNASw supports the following MIBs:

- RFC 2155 APPN MIB with BX Extensions
- RFC 2232 DLUR MIB
- AIW Standard APPN Trap MIB

For further detail and how to use MIBs, see www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml.

In addition, RFC 2353 APPN/HPR in IP Networks and AIW Standard BX and EE are supported.

Note: To ensure the most up-to-date information, be sure that you check the latest software requirement with the Cisco TAC and the IBM support center.

Frequently Asked Questions about APPN-to-SNASw Migration

General

Q. Why do I need to migrate from APPN to SNASw?

A. The traditional APPN feature set, PSNA, will no longer be supported in Cisco IOS Release 12.1 and later. The feature set has been replaced by the new second-generation Cisco APPN product, SNASw. The following milestones have been set for this process:

- EOE for Cisco IOS Release 11.2—April 16, 2001
- End of sales (EOS) and EOE for Cisco IOS Release 12.0—After March 2002

It is time for you and your customer to start considering migration steps from the traditional APPN features to SNASw.

Q. What requirements are addressed by SNASw?

A. SNASw addresses the following requirements:

- To provide the needed SNA routing functionality
- To reduce the complexity in traditional APPN environments by simplifying network designs and reducing configuration requirements
- To address scalability issues associated with traditional APPN NN topologies

Q. What functionality is provided by SNASw?

A. SNASw supports APPN BX, APPN EE, and DLUR functionality.

Q. What version and release of IBM OS/390 is needed to support SNASw BX?

A. To support the BX feature, OS/390 must be running at V2R5 or higher. In addition, the following APARs must be applied: OW37548, OW37549, and OW39559. OW39559 is required for connection network support.

Q. What version and release of IBM CS/390 is needed to support SNASw EE?

A. The *minimum* version and release of CS/390 for EE support is V2R6 with APAR OW36113 applied (the EE function was not present in the initial shipment of Release 6 but was enabled by this PTF). In addition to EE support, APAR OW36113 also provided support for the enhanced Responsive Mode adaptive rate-based (ARB-2) flow control algorithm for HPR over IP. Responsive Mode ARB was released with EE to allow HPR RTP to better compete with TCP for available bandwidth over IP backbones. SNASw supports both the enhanced Responsive Mode ARB (ARB-2) as well as the original ARB (ARB-1) flow control algorithm.

- Q.** What release of Cisco IOS software is needed to support SNASw?
- A.** Cisco IOS Release 12.1 or higher is required.
- Q.** What Cisco router platforms does SNASw run on?
- A.** SNASw is supported on the Cisco 2500, 2600, 3600, 4000, 7200, and 7500 Series routers, as well as the Catalyst 5000 RSM with recommended DRAM (see Cisco CCO Software Center for more information on SNASw DRAM requirements).

Migration Considerations

This section highlights different technical aspects that customers need to consider before starting the migration process.

- Q.** How many upstream APPN links does each of your current NNs support?
- A.** SNASw supports approximately 10-12 host uplinks. If you are planning to have a higher number of links, then you should plan on using SNASw connection network support. SNASw support for connection network makes it possible to reduce the number of predefined links to EN application hosts in your network and significantly reduces configuration effort. A connection network is a single VRN that represents an SATF, which provides any-to-any connectivity for nodes that are attached to it. The VRN is not a physical node; it simply provides a means of defining EN link attachments to other ENs without having to explicitly define links to other ENs with which it communicates.
- Q.** Do you plan to have EE (HPR-only) connections adjacent to interchange transmission groups? Are any of your hosts connected to ICNs?
- A.** A common scenario for this is where an OS/CS/390 EE host is also an SNI gateway to another network (ICN). Prior to OS/CS/390 V2R10 (and releases prior to IBM APAR OW44611), this did not allow sessions to cross-domain subarea partners to exit an ICN via an HPR connection unless the connection was only one hop away from the target EN.

Even with OS/CS/390 V2R10 (or higher) and the fix to IBM APAR OW44611 applied, there will still be one scenario that will not support HPR on an APPN link immediately adjacent to an ICN TG. This is when an ICN defines a connection to a connection network (VRN) and a session is attempted from the subarea network through the ICN and then into APPN over the VRN. The solution to addressing this limitation is to explicitly define the APPN link from SNASw to the ICN. Refer to IBM APAR OW44611 for more information regarding this limitation.

- Q.** Are you currently running APPN NN transport over DLSw+ to remote branch offices?
- A.** The primary goal of migrating from APPN NN to SNASw is the reduction of the total number of APPN NN routers. The two basic approaches to doing this are the following:
 - Replace DLSw+ in every branch with SNASw EE and HPR/IP (the preferred method)
 - Leave the existing DLSw+ network in place to the remote branches, deploy SNASw BX/DLUR in data center (or aggregation layer) DLSw+ peer termination routers to provide emulated NN services (BX) and DLUR support for downstream EN, LEN node, and PU 2.0 devices, and implement SNASw EE uplinks for IP Layer 3 connectivity to S/390 and zSeries NN server/DLUR primary and backup hosts (use SNASw connection network to support dynamic links to other application EN hosts)
- Q.** Do you have all recommended VTAM maintenance for HPR and EE?
- A.** The following information APAR list is the IBM recommended and essential maintenance:

II10953 VTAM HPR RECOMMENDED MAINTENANCE D
NET,VTAMSTOR,MODULE=ISTA UCLA should show UW66546
ENTERPRISE EXTENDER GENERAL INFORMATION

II12223

The following IBM APARs are recommended for customers planning to implement EE:

- APAR OW36113 (OS/CS/390 V2R6)
- APAR OW36458 (OS/CS/390 V2R7)

OS/CS/390 V2R10 (or V2R8 and higher with APAR OW43814 PQ38173) has an enhancement that enables the activation of EE major node definitions prior to TCP/IP startup.

Q. Are you currently implementing APPN RFC 1483 (with or without connection network) over Asynchronous Transfer Mode (ATM)?

A. SNASw provides DLC support compatible with having direct links between SNASw and upstream hosts and, therefore, does not support RFC 1483. When deploying SNASw over ATM, your connectivity choices are LAN emulation (LANE) or HPR/IP EE.

Q. Do you currently have cascaded APPN NN routers supporting DLUR functionality?

A. An APPN architectural restriction of DLUR/DLUS is that a DLUR can never be cascaded below another DLUR. The SNASw DLUR router must be directly connected to the upstream NN/DLUS server.

Q. Is your current APPN network topology composed of multiple APPN NNs cascaded below other APPN NNs?

A. Introducing SNASw BX/EE routers requires a careful insertion strategy if there are currently a number of cascaded NNs in the network path from the downstream device to upstream hosts. SNASw BX/EE must not have traditional NNs below it in the network configuration, so you should add Cisco SNASw routers at the bottom level (layer) of the APPN network topology as a starting point. Direct links or connection network preferably should be used for links between SNASw BX/EE routers and the CS/390 hosts instead of having a complex cascaded NN network in the middle.

This is accomplished using SNASw EE (HPR/IP) such that an IP network can connect the SNASw router EE RTP endpoint directly to upstream hosts. SNASw BX support eliminates the need for any intermediate APPN NN routing by providing all required emulated NN services to downstream SNA devices.

As previously mentioned, when a pre-existing DLSw+ WAN network is in place, SNASw can be deployed at the hub end (data center) or in the distribution (or aggregation) layers such that HPR/IP (EE) transport is supported between the SNASw data center router and the upstream S/390 or zSeries host. At the same time you can support the existing DLSw+ WAN transport of remote SNA traffic downstream into SNASw to use the SNA routing capabilities provided by SNASw BX and support for SNA dependent PU 2.0 devices by SNASw DLUR support.

Q. Do you have any secondary LU-initiated independent LU sessions?

A. Secondary LU-initiated independent LU sessions are not supported by SNASw. SNASw only supports primary LU-initiated independent LU SNA sessions and does not support the session services extensions (SSE) APPN option set.

Q. Does SNASw BX support host-to-host connection between VTAM hosts (APPN border node) or connections to downstream peripheral border nodes?

A. APPN ENs downstream from an SNASw BX router must be real ENs or other branch extenders acting as ENs, not border nodes or peripheral border nodes presenting an EN image. SNASw BX does not support being on the same path between two VTAM hosts (border node or extended border node implementations) or allow AS/400 border nodes acting as ENs to be downstream from a SNASw BX router (as mentioned previously, SNASw does not provide SSE NN server support).

SNASw Performance

Overview

Cisco Systems performed a series of tests to determine the percentage of the CPU utilized on various Cisco router platforms deploying the SNASw feature. This data can help customers make a ballpark comparison between the processing capabilities of different router model types and a determination of how many transactions per second (tps) a particular router platform can support.

The SNASw performance information is divided into three separate categories. First, it provides an approximate benchmark for a new SNASw deployment. Graphs show the impact of data traffic on CPU utilization for a number of Cisco hardware platforms. Second, it shows the performance characteristics of running SNASw alone on the router versus running SNASw and DLSw+ combined. Finally, the data provides a relative performance comparison between SNASw and Cisco's previous APPN PSNA platform.

This performance paper does not compare performance results between SNASw and DLSw+. The decision to deploy SNASw, DLSw+, or a combination of SNASw and DLSw+ should be based on requirements, current environment, and needed functionality. A *DLSw+ TCP Performance* white paper comparing performance results on a wide variety of Cisco hardware platforms running DLSw+ is available at www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/tech/dstcp_wp.htm.

Note: Although processor utilization on most of the router platforms tested was driven to 70 percent router CPU utilization (and higher), it is recommended that networks deploying SNASw and DLSw+ utilize no more than 50 percent of the router CPU. This recommendation is especially important if availability, redundancy, and the ability to scale to meet future growth requirements is an absolute necessity.

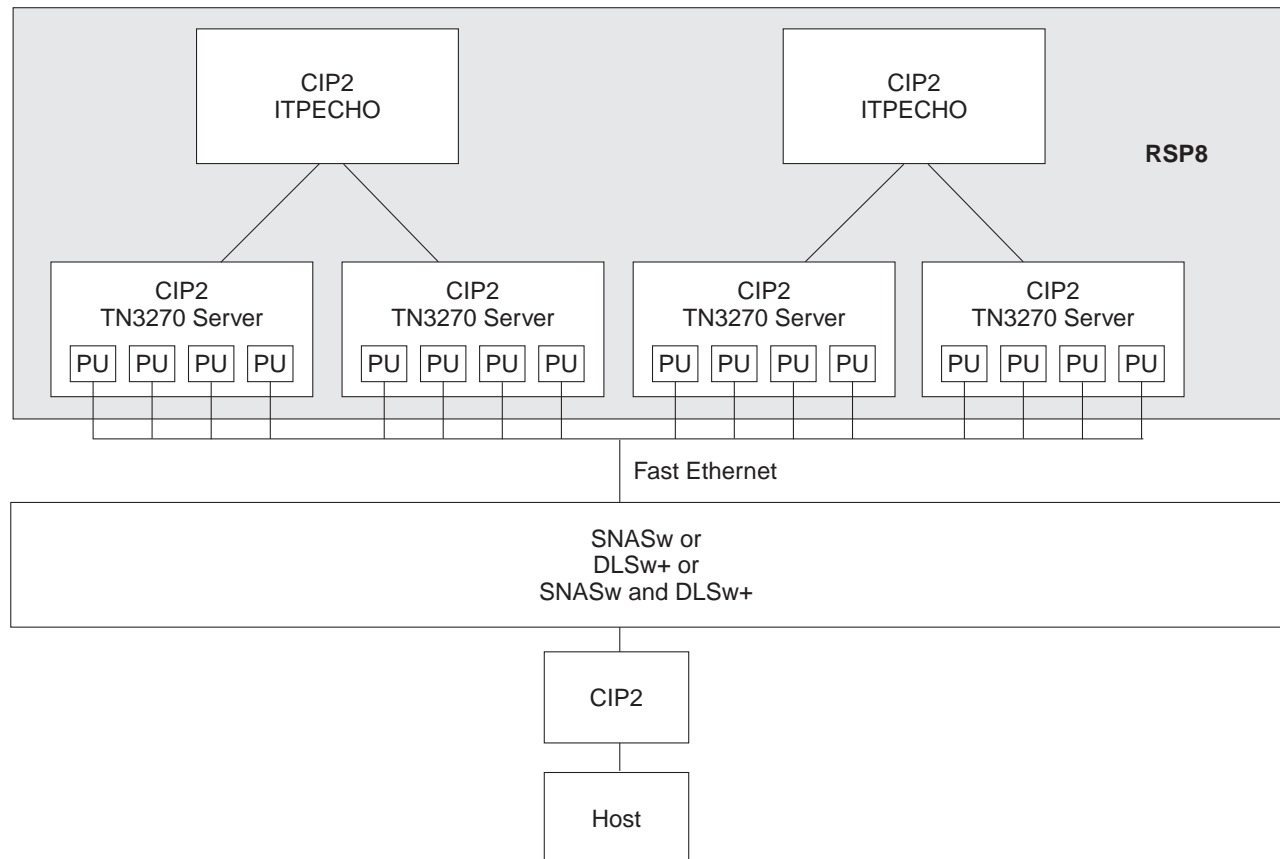
For more information about designing SNASw networks, consult the *SNASw Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/snasw/tech/snasw_rg.pdf. For more information about designing DLSw+ networks, consult the *DLSw+ Design and Implementation Guide*.

The Test Environment

In all performance test cases the ITPECHO test tool was used to generate SNA traffic to an echo application on the host mainframe (ITPECHO is an internal Cisco test tool that runs on a special microcode version for the CIP). Two CIP2 routers were used to run the ITPECHO application, each directed to two Cisco TN3270 Servers for the purpose of generating and directing traffic to SNASw. All six CIPs resided on one Cisco 7513 router.

The traffic generation setup is summarized in Figure F-1.

Figure F-1 Traffic Generation Setup



For each SNA LU created, 100-byte frames were sent to the host every one second with a 1000-byte response frame returned from the host (this transaction profile simulated typical SNA interactive data traffic).

A script was used to start the instances of ITPECHO and record the average CPU usage for a specific transaction per second (tps) rate over five-minute intervals. The script was then restarted using the ITPECHO tool with more LUs (Cisco testing has shown that the load on router CPU is dependent only on the total number of sessions—that is, the CPU load for 100 PUs with one LU per PU and for one PU with 100 LUs per PU is the same), and statistics were recorded for each run. The test procedure was repeated until the CPU on each SNASw router tested was fully utilized (CPU utilization at 100 percent).

Statistics were compiled and recorded, and five-second CPU utilization statistics were recorded in 30-second intervals while data traffic was running. A Cisco IOS `show snasw statistics` command was executed before stopping data traffic. The data frame rate was computed in two ways. The data frames sent and received by each PU were recorded from output of the `show snasw link xid` command before the data traffic was started. Data traffic was terminated after five minutes and the output from the `show snasw link xid` command was recorded again (Cisco internal testing has shown that five minutes is sufficient time to compute an accurate average).


Results

The tests measured the percentage of the CPU utilized on various Cisco routers as a function of transaction rate (number of transactions per second) for SNASw running standalone, as well as SNASw and DLSw+ running concurrently on the same router. In some instances, a variety of modes were examined, including HPR/IP and ISR. The results presented in this paper can assist in making ballpark comparisons between performance characteristics of different Cisco router models and in estimating how many transactions per second a particular Cisco router platform can support.

Branch Router HPR/IP Performance

In many customer environments, SNASw is typically deployed in Cisco routers installed in remote branch offices. This is because the SNASw EE feature can enable the transport of SNA traffic over IP/UDP all the way from the remote branch router into the IBM S/390 or zSeries enterprise mainframe server. Figure F-2 compares the number of transactions per second processed with the corresponding router CPU utilization for the Cisco 2621, 2651, 3620, 3640, 3660, NPE-200, and 4700 hardware platforms in remote branch offices running SNASw EE (HPR/IP).

Figure F-2 Branch Routers Running HPR/IP




Note: The figures in this document represent a linear fit of the average CPU data points.

Data Center Routers Running HPR/IP with DLSw+

Customer enterprises already running DLSw+ for SNA transport over an IP WAN network might opt not to deploy SNASw EE out to remote branches because the DLSw+ network has been in place and stable for a long time. Many of these customers, however, add SNASw to central site routers terminating existing DLSw+ peer connections from remote branch DLSw+ routers. The BX capability of SNASw provides the necessary SNA routing for downstream SNA devices, while the SNASw EE feature transports SNA traffic natively into upstream IBM EE-enabled enterprise servers over IP Layer 3 switches such as the Catalyst 6500 Series. Figure F-3 compares the number of transactions per second processed with the corresponding router CPU utilization for the Cisco RSP2, RSP4, RSP8, NPE-200, and NPE-400 hardware platforms in the data center running SNASw EE and DLSw+ concurrently in the same router.


Figure F-3 Data Center Routers Running HPR/IP with DLSw+



Data Center Routers Running HPR/IP without DLSw+

SNASw can also be deployed in central site data center routers separate from the routers supporting existing downstream DLSw+ SNA transport functions. Figure F-4 compares the number of transactions per second processed with the corresponding router CPU utilization for the Cisco RSP4, RSP8, NPE-200, and NPE-400 hardware platforms in the data center running SNASw EE (HPR/IP) without DLSw+.


Figure F-4 Data Center Routers Running HPR/IP without DLSw+



Comparison of SNASw Modes for Data Center Routers Running with DLSw+

Figure F-5 compares the number of transactions per second processed with the corresponding router CPU utilization running the various supported SNASw modes of operation (EE HPR/IP, HPR over LLC, and ISR) for Cisco RSP2, RSP4, RSP8, NPE-300, and NPE-400 hardware platforms.

Figure F-5 Comparison of SNASw Modes for Data Center Routers Running with DLSw+




Note: The NPE-400 and RSP8 were tested only with SNASw enabled for EE (HPR/IP) and DLSw+.

Comparison of SNASw and PSNA for Branch Router

It is important for customers migrating to SNASw from Cisco's previous APPN NN feature in the Cisco IOS Software (PSNA) to understand the performance differences between the two platforms. Figure F-6 compares the number of transactions per second processed with the corresponding router CPU utilization for the various modes of APPN PSNA and SNASw operation on a representative Cisco remote branch router platform (the Cisco 4700 Series).

Note: No performance results are available for EE with APPN PSNA because the EE feature is not supported on this platform.

Figure F-6 Comparison of SNASw and PSNA in a Cisco 4700 Series Branch Router




Comparison of SNASw and PSNA for RSP4

As in the previous test, it is important for customers to understand the performance differences between APPN PSNA and SNASw running in central site data center environments. Figure F-7 compares the number of transactions per second processed with the corresponding router CPU utilization for the various modes of APPN PSNA and SNASw operation on a Cisco RSP4 hardware platform also running DLSw+.

Note: Again, no performance results are available for EE with APPN PSNA because the EE feature is not supported on this platform.

Figure F-7 Comparison of SNASw and PSNA on the RSP4



Version Information

All routers tested ran the SNASw feature set in Cisco IOS Release 12.1(5), which was the latest SNASw release available at time of testing.

References and Recommended Reading

SNASw Web Site

(www.cisco.com/warp/public/cc/pd/ibsw/snasw/)

“APPN-to-SNA Switching Services Migration”

(www.cisco.com/warp/public/cc/pd/ibsw/snasw/prodlit/swsem_qp.htm)

“Why Migrate to SNA Switching Services”

(www.cisco.com/warp/public/cc/pd/ibsw/snasw/prodlit/snass_pg.htm)

Cisco IOS Bridging and IBM Networking Configuration Guide, Release 12.1, “Configuring SNA Switching Services” (Cisco Documentation CD-ROM or www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ibm_c/bcprt2/bcdsnasw.htm)

Cisco IOS Bridging and IBM Networking Configuration Guide, Release 12.1, “SNA Switching Services Commands” (Cisco Documentation CD-ROM or www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/ibm_r2/br2prt1/br2dsnaw.htm)

DLSw+ Design and Implementation Guide

(www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/prodlit/dlswa_rg.pdf or
www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/prodlit/toc_rg.htm)

SNA Internetworking Design and Implementation Guide

(www.cisco.com/warp/public/cc/so/neso/ibso/data/datacent/datat_rg.htm)

TN3270 Design and Implementation Guide

(www.cisco.com/warp/public/cc/pd/ibsw/tn3270sr/tech/tndg_toc.htm)

IBM APPN Architecture Reference (SC30-3422-04)

IBM APPN Branch Extender Architecture Reference (SV40-0129)

IBM APPN Dependent LU Requester Architecture Reference (SV40-1010)

IBM APPN High Performance Routing Architecture Reference Version 4.0 (SV40-1018)

Inside APPN—The Essential Guide to the Next-Generation SNA (IBM Redbook SG24-3669)

Migrating Subarea Networks to an IP Infrastructure Using Enterprise Extender (IBM Redbook SG24-5957)

OS/390 IBM CS SNA Network Implementation Guide (SC31-8563)

Subarea to APPN Migration: HPR and DLUR Implementation (IBM Redbook SG24-5204)

