Reference Guide

SNA Internetworking Design and Implementation Guide

Volume 3 in the Cisco Internetworking Design Guide Series

Cisco Systems

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

Contents

TABLES

CHAPTER 1

Introducing the Cisco Mainframe Channel Connection

The Cisco Mainframe Channel Connection (CMCC) family of products includes the Channel Interface Processor (CIP) for the Cisco 7000 and 7500 Series routers and the Channel Port Adapter (CPA) line for the Cisco 7200 Series routers. This chapter helps you understand the CMCC solution for today's data centers. The CMCC products support high-performance mainframe access. CMCCs are part of the Cisco Systems solutions to integrate your mainframe with the rest of your network as it evolves to support higher-bandwidth Internet and intranet solutions.

To set the stage for the CMCC family of products, this chapter describes the evolution that is occurring in enterprise networks and data centers. A description of Cisco's history and expertise in mainframe channel technologies and then an overview of the CMCC family follow this. The chapter concludes with a detailed description of the CIP and the CPA line and a summary of product capabilities and differences. For more information about CMCCs, see www.cisco.com/warp/public/779/largeent/sna/edc/mcc.html.

Evolution of the Enterprise Network

In the 1980s, the term SNA network was synonymous with enterprise network. Developed by IBM to support the computing and networking needs of large enterprises, SNA dominated as the networking architecture of choice for most large enterprises. SNA delivered high levels of scalability and availability to enterprises that were unsurpassed by any other networking architecture of the time. SNA was a hierarchical architecture with the mainframe at the pinnacle, which reflected computing realities of the 1970s and 1980s. Mainframes, residing in data centers, were the repository of the majority of mission-critical applications. End users accessed mainframe applications from teletype machines or display terminals (known as 3270 devices). IBM developed SNA to define how display terminals could access applications and information in IBM S/390 mainframes. It should be noted that, throughout the remainder of this document, the term *IBM S/390* refers to both the traditional System/390 mainframe and the new IBM z900 processors. Similarly, the term *OS/390* refers both OS/390 and its successor, zOS.

Figure 1-1 shows a simple SNA network and some of its key components.

Figure 1-1   Simple SNA Network

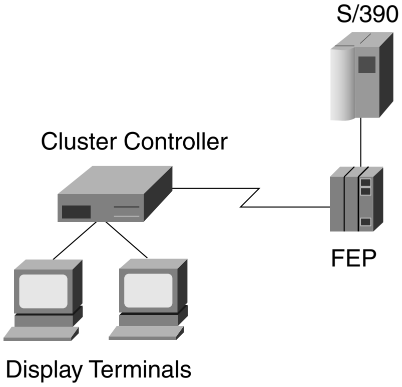

Beginning in the 1980s, many advanced enterprises began to integrate TCP/IP into their networks to accommodate interoperability between a wide variety of systems, including UNIX-based servers. Today, virtually all enterprises have IP in at least some portion of the network to support corporate intranets and Internet access. However, the presence of IP does not mean that SNA has been eliminated. Many enterprises continue to support SNA applications, SNA devices, or SNA networking protocols.

There are a variety of ways that an enterprise can accommodate both SNA and IP in the enterprise. First, it can keep the two environments completely separate. This eliminates any issues of integrating the two environments but can be very costly. Second, it can build a common backbone, usually based on IP protocols, and keep the end systems and applications either native SNA or native IP. This approach helps to eliminate the costs associated with supporting different networks but means that some desktops are supporting two different protocol stacks—SNA and IP. In the third approach, the SNA stacks are eliminated from the desktops and SNA applications are accessed by utilizing a standard that allows IP clients to access SNA applications (TN3270), which can dramatically reduce the costs associated with desktop maintenance and support. In the final approach, the enterprise converts, over time, its SNA applications to IP-based equivalents. Because today's mainframes come equipped with an IP stack, this final step does not mean that an enterprise must eliminate its mainframes.

Figure 1-2 shows the four approaches to SNA-to-IP migration. Cisco offers a wide variety of solutions, including its CMCC solutions, to support an enterprise at each stage in the migration. Many enterprises will find, in fact, that their networks have elements of each of the four phases. Cisco solutions offer the ability to independently migrate portions of the network based on business need.

Figure 1-2 SNA-to-IP Migration

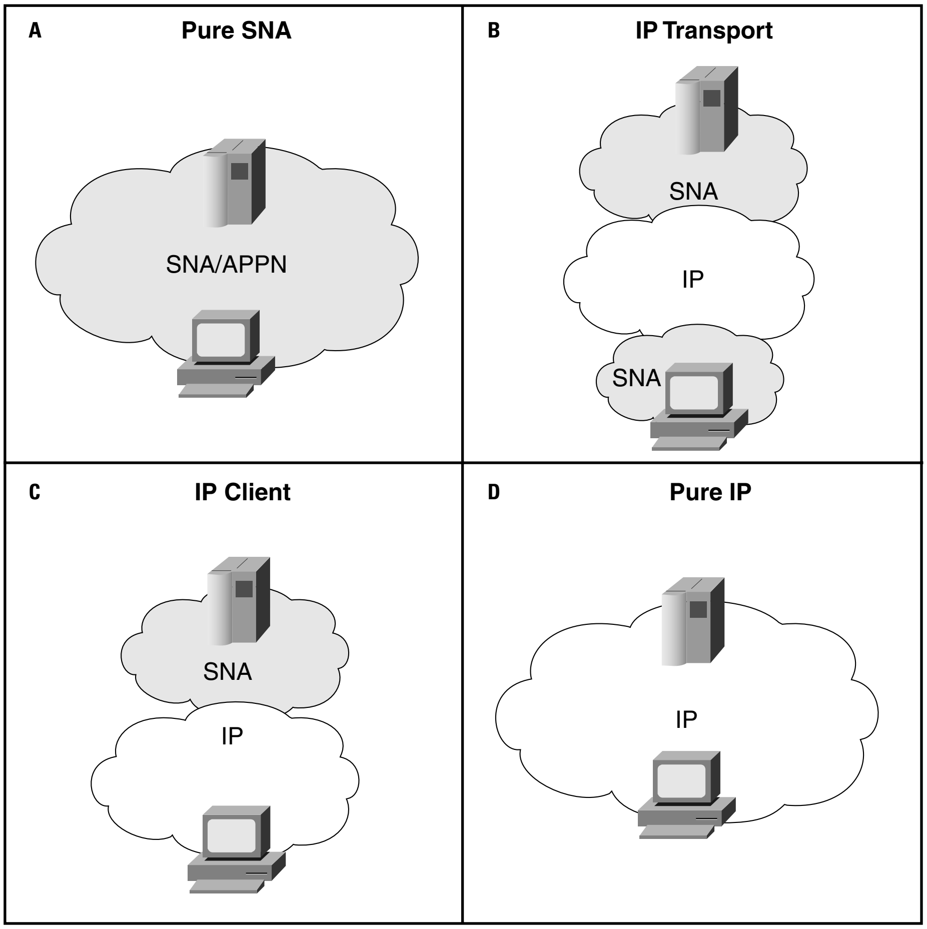

Quadrant A: Pure SNA

Many IBM networks today still access SNA applications on the mainframe from SNA clients. In Figure 1-2, Quadrant A represents a pure SNA network that is fully based on the IBM protocols and architecture. All host and client systems are SNA-based. Cisco supports a Quadrant A environment through the use of the Cisco SNA (CSNA) feature on the CMCC. SNA traffic from one or several serial lines or Token Ring or Ethernet LANs can be aggregated in a single data center router, such as a Cisco 7500 or 7200 Series router. Using CSNA, this SNA traffic can be sent to the mainframe via either Enterprise Systems Connection (ESCON) or parallel bus and tag connections.

Quadrant B: IP Transport

More than 87 percent of enterprise networks have migrated their backbones to TCP/IP. Quadrant B represents the first major step in the migration from SNA to IP. In Quadrants A and B, the branch and data center fundamental network characteristics have not changed. However, the backbone network in Quadrant B has been replaced with IP.

Cisco supports a Quadrant B environment using technologies such as Data-Link Switching Plus (DLSw+), remote source-route bridging (RSRB), and Advanced Peer-to-Peer Networking (APPN). These technologies run in the Cisco router and connect to SNA applications in the mainframe through a CMCC using the CSNA feature. For more information on DLSw+, refer to the *DLSw+ Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/prodlit/dlswa_rg.pdf. For more information on the SNA Switching Services (SNASw) feature, which provides support for APPN nodes, refer to the *SNA Switching Services Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/snasw/tech/snasw_rg.pdf.

Quadrant C: IP Client

With the proliferation of Internet connections and the fact that TCP/IP is included free with Windows 95/98 and Windows 2000, more organizations are looking at TN3270 as a low-cost means to access some of their SNA applications. TN3270 eliminates the requirement for dual stacks on the desktop and minimizes the cost of specialized desktop software. In Quadrant C, SNA is isolated to the data center, and the desktop has TCP/IP only.

Cisco supports a Quadrant C environment through gateway functions that are provided in the router and CMCC. Two examples are the Cisco Transaction Connection (CTRC), which provides access to S/390-based Customer Information Control System (CICS), Information Management System (IMS), and DB2 data, and the Cisco TN3270 Server application, which runs on a CMCC. For more information on the TN3270 Server, refer to the *TN3270 Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/tn3270sr/tech/tndg_toc.htm.

Quadrant D: Pure IP

Finally, in Quadrant D, many organizations are building new mainframe applications or rewriting existing applications to use TCP/IP.

Cisco supports a Quadrant D environment through the use of IP Datagram support and the support of SNA traffic through APPN over IP (APPN/High Performance Routing [HPR] over IP) using the SNASw Enterprise Extender (EE) feature. Connectivity to the S/390 can be through an ESCON or parallel channel, through the use of a CMCC, or through an Open Systems Adapter (OSA). For more information on the SNASw feature, which provides support for APPN nodes, refer to the *SNA Switching Services Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/snasw/tech/snasw_rg.pdf.

Evolution of the Data Center

The previous section discussed trends that have been impacting enterprise networks over the last two decades. These larger trends, and the migration from SNA to IP, have impacted the entire network, from the data center to the remote campus and even to the smallest branch office. There are, however, trends that are specific to the evolving nature of the typical enterprise data center that are relevant to the discussion of the Cisco CMCC product family, such as:

- Changing role of the mainframe
- Centralization of servers
- Increased sophistication and bandwidth of the data center backbone

Changing Role of the Mainframe

In the 1970s and early 1980s, the mainframe was the linchpin of the enterprise, running the majority of all applications, particularly mission-critical applications. With the advent of the client/server revolution beginning in the mid-1980s, some industry experts predicted the decline and eventual replacement of the mainframe by UNIX and PC servers that would run new applications based on the client/server paradigm. That prediction turned out to be highly optimistic.

Although it is true that UNIX and PC servers offered, at the time, an attractive price/performance ratio, these servers did not offer the same level of fault tolerance and disaster recovery capabilities of the mainframe. Many IT organizations could not justify rewriting their mission-critical applications on client/server platforms. And in the last decade, dramatic and rapid improvements have been made to mainframe technology. Today's mainframe offers an exceptional price/performance ratio. As a result, the mainframe remains a staple in most large enterprises.

The role of the mainframe, however, has changed. It is still accessed directly by end users via software that emulates a traditional terminal display, although to a lesser extent than in the past. Now, to an increasing extent, the mainframe acts as a repository of enterprise data to a growing and sophisticated set of servers such as Web servers, Web application servers, and host integration servers. Another shift is the pervasive use of TCP/IP on the mainframe to support new applications. These shifts require that today's mainframe have very efficient and high-speed access to the enterprise network.

Centralization of Servers

At the beginning of the client/server revolution, most servers were placed out in the network, near the department that the application on the server was supporting. Eventually, most large enterprises discovered that supporting a large base of far-flung servers, each potentially running a different type or level of operating system, was a very expensive proposition. As a result, most large enterprises today have centralized at least some types of servers back into the data center. Server farms, created with a large number of high-end servers rack-mounted and resident in the data center, are now the norm.

The implication of the server centralization to the mainframe and the data center is profound. In the 1980s, the data center contained few LANs and many serial lines that connected the data center to the remote enterprise locations. With the reconcentration of computing power and traffic back into the data center, the data center network is now extremely sophisticated and supports very high bandwidth. There is a very large amount of traffic that flows between the different servers in the server farm, as well as between the server farm and the mainframes.

Increased Sophistication and Bandwidth of the Data Center Network

The modern data center network is built with a high-speed, switched LAN infrastructure. Often running at speeds of 1 Gbps or faster, the data center backbone has the bandwidth in place to support huge amounts of traffic. The servers in the data center often have dedicated, high-speed LAN connections. This is in direct contrast with the traditional SNA environment, in which a "high-speed link" was a serial line running at 56 KBps and Token Ring LANs started at a paltry 4 Mbps. The mainframe, an integral part of the data center computing environment, must be able to participate in this high-speed environment. The connectivity choices of a decade ago, such as the front-end processor (FEP) and the IBM 3172 Interconnect Controller, are not able to address these requirements.

In addition to the high bandwidth supported, the data center network is a very sophisticated blend of technologies. With the prevalence of corporate intranets and Internet connectivity, the data center supports a wide variety of security devices, directory services, policy servers, load balancing devices, and so on. The mainframe does not necessarily participate directly with these networking appliances and devices, but the devices that connect the mainframe to the network must support these emerging network services.

Cisco Mainframe Channel Connectivity

IBM channel-attachment support is provided on the Cisco 7000 and 7500 Series routers by the CIP and on the Cisco 7200 Series routers by the CPA. With a CMCC, the Cisco router can directly connect the mainframe to the internetwork in the data center. Cisco was a pioneer in directly connecting the modern router to the mainframe and has been offering direct mainframe channel connectivity since 1996. Its customer base is worldwide, with thousands of enterprises using the CMCC family.

The CMCC adapters are based on IBM technology and support ESCON, ESCON Director, and bus and tag block multiplexor channel connections. The CMCC supports both SNA and TCP/IP traffic to and from the mainframe and can replace or augment traditional mainframe connectivity devices such as the FEP and the IBM 3172 Interconnect Controller.

An ESCON Director greatly reduces the number of channel adapters and physical cable connections required to share devices among multiple systems. One control unit connection to the ESCON Director can provide all of the required connectivity for a multi-image configuration. It also can handle multiple concurrent data transfers. If the ESCON Director is configured with 60 ports, 30 pairs of ports can transfer data at channel speeds. A router with a CMCC can connect to multiple hosts, using the ESCON Director with a single ESCON interface.

The CIP and the CPA support the same CMCC software applications. The differences between the CIP and CPA are of performance and capacity. The difference in performance is based on the internal bus architecture of the CIP and the CPA. The difference in capacity is based on maximum memory configurations (128 MB for the CIP and 32 to 128 MB for the CPA).

Channel Interface Processor

The CIP is a channel-attached interface for the Cisco 7000 and 7500 Series routers. The CIP connects a host mainframe to a control unit and, in many cases, eliminates the need for a FEP or interconnect controller for channel-attachment support. The CIP is designed for high-end network environments that demand high-performance, high-port density, and high-capacity solutions. The CIP supports the IBM ESCON Channel Adapter (ECA) and bus and tag Parallel Channel Adapter (PCA) channel-attached interfaces from the Cisco 7000 and 7500 Series routers to IBM mainframes.

The CIP offers the following benefits:

- Maximum throughput for every application
- Maximum port density
- High-speed processing engines

A single CIP can support up to two physical channel interfaces of either PCA or ECA in any combination. The CIP parallel channel interface is provided by the PCA, while the ESCON channel interface is provided by the ECA. Each CIP is configured with the appropriate channel adapters at the time of manufacturing.

The Cisco 7000 and 7500 Series routers support online insertion and removal (OIR), which allows you to install or remove CIPs while the system is operating.

Channel Port Adapter

The CPA is available for the Cisco 7200 Series routers. The CPA expands the value of the Cisco IBM channel solution by providing channel connectivity to midrange mainframe configurations. The CPA is a standard single-width port adapter supporting ESCON or parallel channel interfaces to IBM mainframes and IBM-compatible mainframes.

The CPA offers the following benefits:

- Support for all major LAN and WAN interfaces
- Superior performance in the midrange market segment

Each CPA provides a single channel interface for the Cisco 7200 Series routers. In some cases, the CPA eliminates the need for a separate FEP or interconnect controller. The CPA contains a single input/output (I/O) connector.

The Cisco 7200 Series router supports OIR, which allows you to install or remove port adapters while the system is operating.

The CPA is available in three forms: the ESCON CPA (ECPA), the high-performance ESCON CPA Version 4 (ECPA4), and the Parallel CPA (PCPA), described in the following sections.

The ECPA is a high-speed port adapter. A single Cisco 7200 VXR Series router can support up to five high-speed port adapters. Each ECPA model is available with 32 MB of system memory and is capable of supporting one ESCON port.

The ECPA4 is a high-performance version of the ECPA. The ECPA4 system processor's performance is more than twice that of the ECPA processor and comes with four times the memory (128 MB). The ECPA4 includes an updated ESCON chip set.

The PCPA provides support for a single parallel channel physical interface, with both 3.0 and 4.5 MBps data transfer rates. The PCPA uses the same processing engine as the ECPA and supports 32 MB of system memory.

Table 1-1 shows the differences among the different types of CMCC adapters.

Table 1-1 Product Differences among the CIP, ECPA, ECPA4, and the PCPA

Feature	CIP	ECPA	ECPA4	PCPA
Router Platform	Cisco 7500 and Cisco 7000 with RSP7000	Cisco 7200	Cisco 7200	Cisco 7200
Channel Interfaces	ESCON, Parallel	ESCON	ESCON	Parallel
Maximum Number of Interfaces	2	1	1	1
Maximum Memory	128 MB	32 MB	128 MB	32 MB
Cisco IOS® Release Support	Cisco IOS Release 10.2 and later	Cisco IOS Release 11.3(3)T and later	Cisco IOS Release 12.1(5)T and later	Cisco IOS Release 11.3(3)T and later
Virtual Port Adapter	2	0	0	0
Channel Interface State Tracking (HSRP, SNMP Alerts)	Yes	*Disabled*—Use the **state-tracks-signal** command to enable	*Disabled*—Use the **state-tracks-signal** command to enable	*Disabled*—Use the **state-tracks-signal** command to enable

CHAPTER 2

Introducing SNA on the CMCC

Although the vast majority of enterprises have TCP/IP within the enterprise network to meet pressing business needs such as offering services via the Internet, that does not mean that SNA has disappeared from enterprise networks. More than 70 percent of Fortune 1000 companies still base some of their mission-critical applications on SNA. More than 750,000 SNA gateways and controllers currently are installed and providing end users with access to SNA applications. Cisco supports end-to-end SNA networks with the Cisco IOS Software, as well as with special-purpose hardware that attaches Cisco routers to the S/390 environment.[1]

This chapter introduces the features that a CMCC offers in an SNA environment. The information in this chapter can help you determine when and where to use the CMCC and IBM services of the Cisco IOS Software, as well as the best design for your data center to optimize performance and availability.

This chapter includes the following information:

- An overview of SNA for readers who are familiar with Cisco routers but not familiar with SNA networking
- An overview of APPN, APPN/Intermediate Session Routing (ISR), and APPN/HPR
- An overview of the SNASw features: SNASw Branch Extender (BX), SNASw Dependent Logical Unit Requester (DLUR) support, and SNASw Enterprise Extender (EE) features
- A description of how SNA and APPN devices can access an SNA mainframe using CMCC and other Cisco SNA features

Overview of SNA

SNA was invented in 1974 as a standard way to access applications on IBM mainframes. SNA has evolved and changed since then, but many networks still run the original architecture, so it is important to understand the following basic SNA concepts:

- Subarea SNA
- APPN
- SNASw

Overview of Subarea SNA

The original SNA architecture was also known as subarea SNA, because SNA networks were divided into logical groupings, called subareas, which facilitated routing and directory functions. Subarea SNA was hierarchical. That is, the applications and network control software resided on IBM mainframes, not on workstations.

[1] Cisco provides a number of different channel-attachment hardware features. These features include the CIP and the CPA, discussed in the chapter Introducing the Cisco Mainframe Channel Connection. The features provided by the CIP and CPA are the same; the difference is the *scalability* of each solution.

The Virtual Telecommunications Access Method (VTAM), which is now part of the IBM Communications Server for OS/390 (CS/390), was the mainframe software that controlled SNA subarea networks. VTAM used a system services control point (SSCP) to establish a control session with dependent SNA devices within its span of control, or domain. VTAM is responsible for the following SNA control functions:

- *Activating and deactivating SNA devices*—Similar to a "logical power-on" of the device
- *Providing directory functions for SNA devices*—Finding the correct mainframe logical partition (LPAR) running the application
- *Assisting in session establishment*—Like a "telephone operator" for SNA communication
- *Routing SNA*—Routing traffic toward the destination SNA application
- *Receiving alerts of events*—Receiving notification of what is happening in the network

In general, all devices within a VTAM's domain of control had to be configured to that VTAM. Later VTAM releases allowed dynamic configuration of resources using generic definitions. A VTAM can also dynamically find resources owned by another VTAM domain; such resources are known as cross-domain resources (CDRSCs).

There are few pure subarea SNA networks still in existence. Enterprises that have continued to invest heavily in SNA equipment have often introduced APPN, the IBM follow-on to subarea SNA. Many others have instead migrated their networks, in part or fully, from subarea SNA to TCP/IP. There are, however, few enterprises at either end of the spectrum—pure SNA/APPN or pure TCP/IP. Most have a blend, and many continue to have a sizable installed base of traditional, subarea SNA equipment. This installed base, although declining over time, must be accommodated and integrated during the migration stages to protect the investments the enterprise has made in SNA applications and infrastructure.

Physical Unit Types

SNA uses the term physical unit (PU) followed by a number (1, 2, 2.1, 4, or 5) to identify network processors that can participate in SNA networks. The number indicates the specific SNA functionality provided by each PU type. Figure 2-1 shows the components of a subarea network, which are described in this section.

Figure 2-1 Subarea SNA Network Components

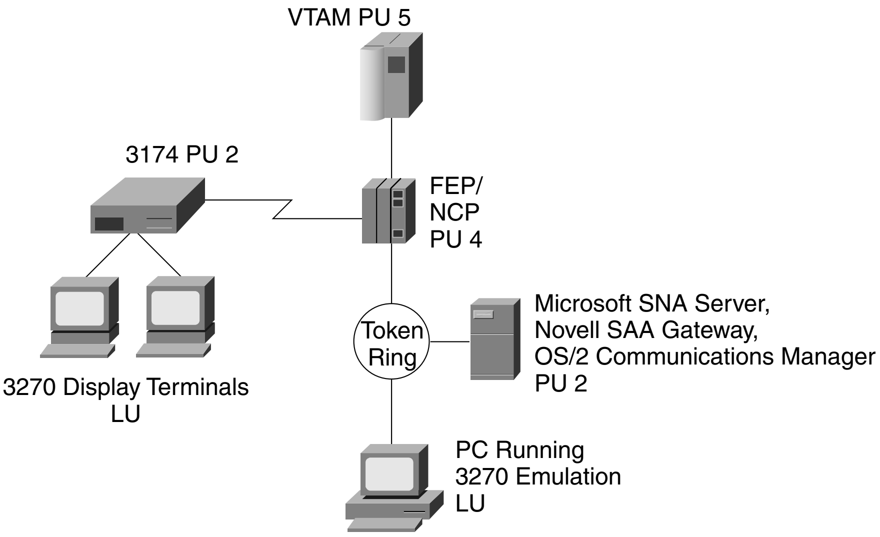

PU 5 Functionality—VTAM

A PU 5 provides subarea SNA routing functionality and is generally implemented in VTAM, along with the SSCP. The SSCP provides connection point and network management services for a specific set of PU 4, PU 2, and PU 1 nodes, known as a domain.

PU 4 Functionality—NCP

To offload some of the mainframe processing, IBM developed FEPs that communicate with the mainframe over communication channels. The FEPs run the Network Control Program (NCP), which routes SNA traffic to the mainframe that runs the destination application. The subarea routes must be statically configured in an NCP, but VTAM dynamically selects the first route that matches the requested SNA Class of Service (COS) and destination subarea number. The NCP also prioritizes traffic on its outbound queues, based on the transmission priority assigned to a particular SNA session. Finally, the NCP provides a boundary function that enables devices on the boundary of the SNA network, such as cluster controllers, to access the mainframes. The NCP implements PU 4 functionality.

PU 2 Functionality—Cluster Controllers and PC Gateways

Cluster controllers (such as IBM 3174s) provide access to SNA networks from display terminals or terminal emulators. Cluster controllers access the SNA network through an SNA boundary node, such as the NCP or VTAM, but they do not provide SNA routing or COS functions. Cluster controllers sometimes are called peripheral devices because they are located on the periphery of an SNA network and do not fully participate in all the SNA functionality. Cluster controllers provide PU 2 functionality.

Special software running on personal computers and other end systems supports 3270 emulation to allow these end systems to communicate with existing 3270 mainframe applications. Server-based PC gateways and Cisco routers running certain SNA features provide the same PU 2 functionality as provided by cluster controllers. Some client software implements both PU and logical unit (LU) functionality.

Logical Unit Types

Display terminals, such as 3270 terminals, enable end users to request application services. End users enter keystrokes, which are sent to the cluster controller. The cluster controller places the keystroke information in an SNA request unit (RU) with a boundary (peripheral) format-identifier 2 (FID2) header and forwards the RU to an NCP. The NCP converts the header from a FID2 to a FID4 (converting local addresses to subarea addresses and adding a transmission priority field) and forwards it to the next hop in the SNA network. Eventually the RU reaches an SNA mainframe application in which the request is processed, and the results are returned to the display terminal. Because application processing is performed on the mainframe, every request and its associated response must travel the network before the response is displayed. Applications, printers, and 3270 terminals or emulators are known as LUs. Several LU types are available in subarea SNA:

- 3270 terminals and emulators appear as LU 2s to a VTAM.
- Printers generally appear as LU 1s or LU 3s to a VTAM.
- LU 0 applications use an unstructured data field to enable advanced functions.
- Advanced Program-to-Program Communications (APPC) applications communicate using LU 6.2, which provides peer-to-peer communication between programs. Unlike display terminals, PCs can run applications, communicating program to program with a mainframe application rather than asking a mainframe application to process a request.

CMC Environment

In a Communication Management Configuration (CMC) environment, a single VTAM (the CMC host) owns all the SNA resources in the network and is involved in the initiation and termination of every session. The other mainframe images support only SNA applications. CMC design keeps the burden of network processing off of the application hosts and simplifies the collection of management data.

Overview of APPN

In the early 1980s, it became apparent that the hierarchical architecture and static definition of subarea SNA were major impediments to supporting new systems within the enterprise that utilized distributed client/server or peer-to-peer technologies. Thus, in 1985 IBM developed APPN to allow SNA devices and applications to participate in peer-to-peer sessions. APPN also provides dynamic routing and dynamic directory capabilities and extends SNA service levels and prioritization farther out in the network.

APPN is not a hierarchical network architecture but a peer architecture. Three major node types can exist in APPN networks:

- *Network nodes (NNs)*—These nodes are SNA routers, responsible for locating resources, selecting paths, and working with the users to set up sessions.
- *End nodes (ENs)*—These nodes are application hosts, end users, or controllers representing multiple users.
- *Low-entry networking (LEN) nodes*—These nodes represent early, pre-1985 APPN technology.

The following list is a summary of APPN characteristics:

- The topology of the network is not predefined. NNs exchange information so that each has an entire picture of the network, all the NNs and the links connecting them. Each NN also maintains a local topology, the ENs and the links between ENs and NNs.
- Directory services are distributed. Each NN knows about the resources attached to its ENs, plus other network resources that have sessions with its resources. Locations of network resources are determined via broadcast.
- SNA COS enables the selected path to deliver an appropriate service level and prioritize messages to ensure that the service level is maintained.
- Support for DLUR/Dependent LU Server (DLUS) provides dependent SNA device support (PU2/LU2) over APPN networks.

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 first looks in its directory to determine 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 broadcasts 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 instructs the EN to complete session setup. Figure 2-2 illustrates an APPN network.

Figure 2-2 Sample Original Architecture APPN Network and Associated Directory and Topology Databases

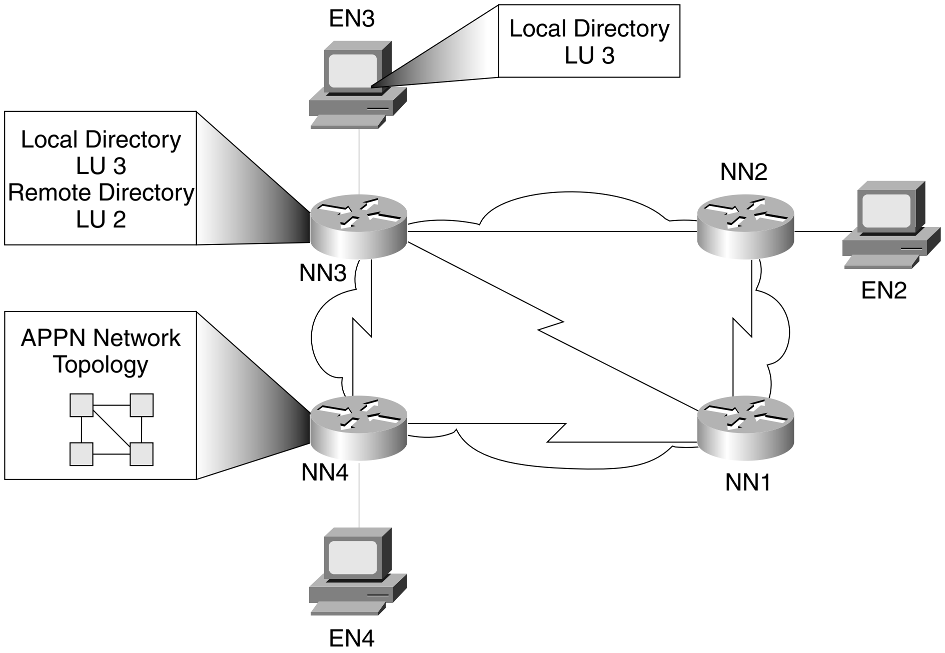

APPN ENs provide local directory services and communicate with their NN server to access other resources in the network. APPN ENs dynamically register their local resources with their upstream NN server.

APPN nodes communicate without the assistance of the mainframe VTAM SSCP. Instead of having a single control point in the mainframe, every EN and NN has its own control point that controls its local resources (applications and links). LEN nodes, which predate APPN support, implement a rudimentary subset of distributed SNA PU 2.1 functionality and require substantial configuration. They are not discussed further in this chapter.

Because APPN is more dynamic and has many more functions than subarea SNA, one might expect it to have quickly overtaken subarea networks. This did not happen for several reasons. First, subarea SNA was not initially supported by APPN. Second, until Release 7.4, the NCP could participate in APPN only when combined with VTAM as a composite network node (CNN). Last, and more important, when APPN was first invented, it supported only APPC LU 6.2 applications. Most mainframe applications were dependent LU applications (3270 or LU 0 applications). VTAM 4.2 addressed this problem with a feature known as DLUS, which is discussed later in this chapter.

APPN has evolved since its original release to overcome these limitations and to support emerging enterprise requirements. It should be noted that some enterprises adopted APPN in its initial stages and became disillusioned by its initial limitations. However, today's APPN technologies (that is, SNASw) offer capabilities that provide superior scalability and support the migration to TCP/IP. Figure 2-3 depicts the steps in the evolution of APPN. These steps are detailed in subsequent sections of this chapter.

Figure 2-3 Evolution of APPN

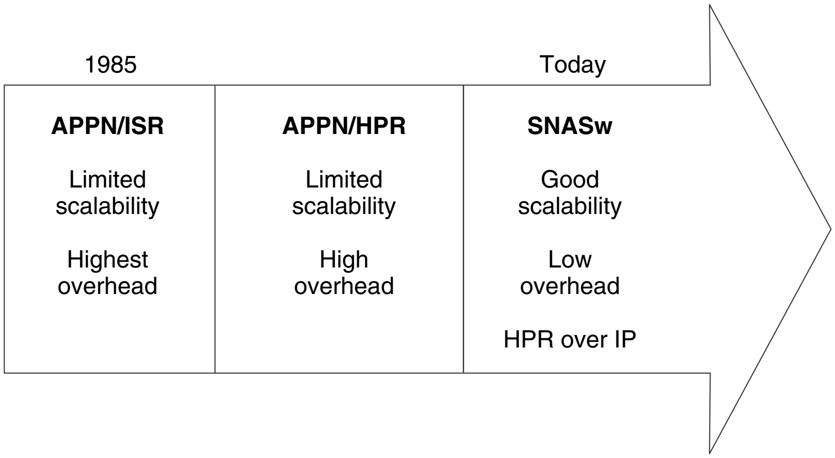

Overview of APPN/ISR

The original (first-generation) APPN architecture utilized APPN NNs to forward SNA session traffic using ISR. ISR provides node-to-node, connection-oriented, data-link control, which provides hop-by-hop error correction and retransmission. The NNs between two APPN endpoints participate in the hop-by-hop guarantee of delivery. This participation causes every APPN data path information unit (PIU) to be examined and processed by the high-level portions of the APPN software in every NN.

This functionality causes too much overhead and limits the scalability of the solution. ISR is processor-intensive. (It is equivalent to running a full TCP stack in every hop along the path.) Also, APPN ISR does not support nondisruptive rerouting around link failures.

Overview of APPN/HPR

In second-generation APPN, HPR provides a connectionless layer for SNA routing called Automatic Network Routing (ANR) with nondisruptive routing of sessions around link failures. HPR also provides a connection-oriented layer called Rapid Transport Protocol (RTP), which supports end-to-end flow control, error control, and sequencing.

Conceptually, RTP is like TCP. RTP is a reliable, connection-oriented protocol that ensures data delivery and manages end-to-end network error and flow control. RTP creates new routes following a network failure. RTP nodes establish RTP connections to carry session data. All traffic for a single session flows over the same RTP-to-RTP connection and is multiplexed with traffic from other sessions using the same connection. The RTP layer is invoked only at the edges of an APPN network. In intermediate nodes, only the ANR layer is invoked. ANR is a connectionless service that is responsible for node-to-node, source-routed service.

Overview of SNASw

HPR resolved some problems of APPN ISR. However, HPR did not address one major problem that limited the scalability of APPN-based networks—the large amount of SNA topology and broadcast search traffic generated by APPN NNs. When an APPN network grew to a certain size, this type of traffic could consume much of the available bandwidth. As a result, few large enterprises adopted APPN throughout their networks.

Cisco SNASw represent an evolution in APPN to address the scalability limitations of earlier APPN technology and also to support the trend in enterprise networks toward IP infrastructure. There are two features that comprise SNASw. The BX feature directly addresses the scalability limitations of earlier APPN technology by effectively reducing the number of NNs within the network. The EE feature transports APPN data over IP/User Datagram Protocol (UDP) transport using the HPR-over-IP capability defined in RFC 2353.

Because many of APPN's shortcomings have been addressed with the BX and the EE features, many enterprises now are considering using APPN for their data centers. The BX and EE features allow you to leverage CMCC deployment by using SNASw BX to replace necessary SNA application routing functionality previously provided by the FEP.

Overview of the SNASw BX Feature

Cisco recommends using the BX feature for any APPN network to reduce topology and locate broadcast traffic. Using BX, SNASw appears like an EN upstream and therefore does not participate in topology updates and locates as APPN NNs do (which allows SNASw networks to scale). It provides a NN image and NN services to downstream SNA devices. SNASw can register downstream devices to the VTAM NN central directory server and provides DLUR function, which is discussed in a later section. Figure 2-4 shows the SNASw BX feature.

Figure 2-4 SNASw BX Feature

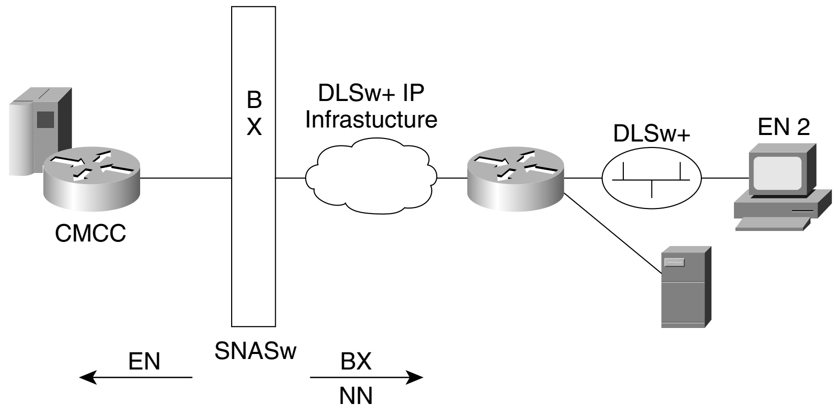

Overview of the SNASw EE Feature

The SNASw EE feature transports SNA data using UDP/IP encapsulation. EE enables transport of HPR data over a native IP network, without requiring DLSw+ transport. The RTP component of HPR provides reliable delivery of frames and flow control.

The SNASw EE feature offers nondisruptive rerouting between the APPN RTP endpoints (the nodes running the EE function). SNASw EE also enables end-to-end, nondisruptive rerouting around links and failures, and it preserves SNA COS end to end. Figure 2-5 shows the SNASw EE feature.

Figure 2-5 SNASw EE Feature

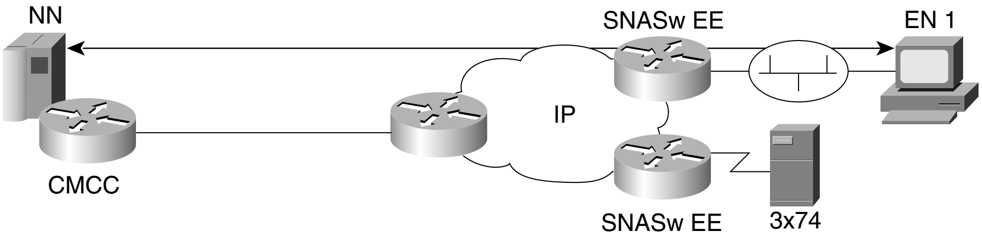

Overview of the SNASw DLUR/DLUS Features

DLUR/DLUS is a feature that was added to APPN to allow SNA subarea traffic to flow on an APPN network. Before this feature, APPN assumed that all nodes in a network could initiate peer-to-peer traffic (for example, sending the BIND to start the session). Subarea SNA end devices, referred to as dependent LUs (DLUs), cannot initiate peer-to-peer traffic and require VTAM running on the mainframe host to notify the application, which then sends the BIND to the end device.

The APPN architecture provides support for subarea DLUs through DLUR/DLUS. DLUR/DLUS allows the control traffic between a VTAM and a subarea DLU to be transported over an APPN network. SNASw supports the DLUR function for SNA dependent devices (PU 2.0) while the VTAM NNs provides DLUS support.

To initiate the subarea SNA sessions, a client/server relationship must exist between APPN DLUS running on the S/390 host and the Cisco SNASw DLUR router, which supports DLUR. A pair of LU 6.2 type sessions is established between the DLUR and DLUS, with one session established by each endpoint. These sessions are used to transport SNA subarea control messages that must flow to activate the DLU resources and initiate their LU-to-LU sessions. Figure 2-6 shows the DLUR feature.

Figure 2-6 SNASw DLUR Feature

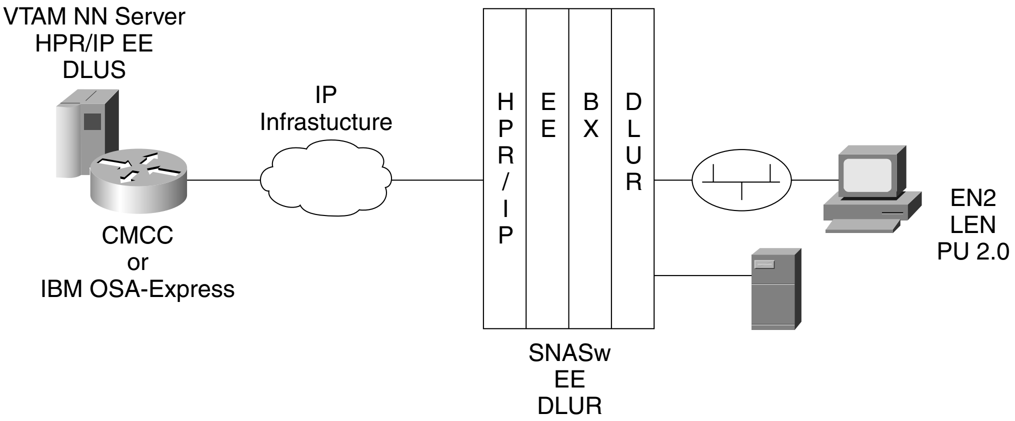

Connecting SNA and APPN Devices Using CMCC

The Cisco channel-attached router provides connectivity between many diverse SNA devices, as shown in Figure 2-7. Using a Cisco 7000, 7200, or 7500 Series router with a CMCC and Cisco SNA (CSNA) support enabled, you can connect two mainframes (either locally or remotely), connect a mainframe to a PU 2.0 or 2.1 device, or connect a mainframe to a FEP in another VTAM domain. (VTAM does not support FEP ownership through an external communication adapter [XCA] device, so a local NCP must activate the remote FEP.)

Figure 2-7 Connectivity among SNA Devices

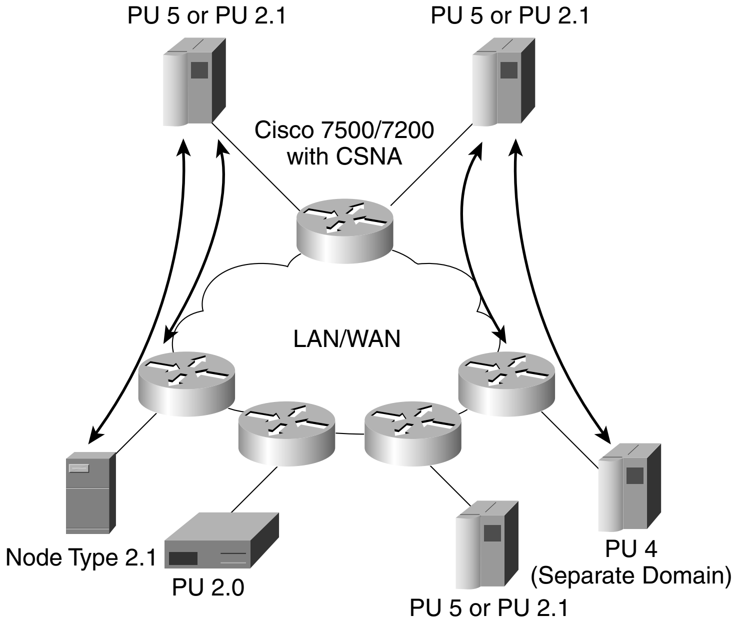

Many options are available for connecting SNA devices (PU 2s or PU 2.1s) to a CMCC-attached router. Access to the CMCC is either using Logical Link Control, type 2 (LLC2) and an internal virtual Token Ring (regardless of the medium the end system is using) or HPR/IP support using SNASw EE. SNA functions, such as DLSw+ or SNASw can reside either in the channel-attached router or in a central campus router that is connected to the channel-attached router.

Local LAN-attached resources or campus resources connected to an Asynchronous Transfer Mode (ATM) LAN Emulation (LANE) backbone can bridge into the CMCC-attached router. Local Synchronous Data Link Control (SDLC) devices can attach directly to a Cisco router and use either DLSw+ local switching or SNASw to access the CMCC via LLC2 or HPR/IP (EE). This example is shown in Figure 2-8.

Figure 2-8 Connecting Local Resources to the CMCC

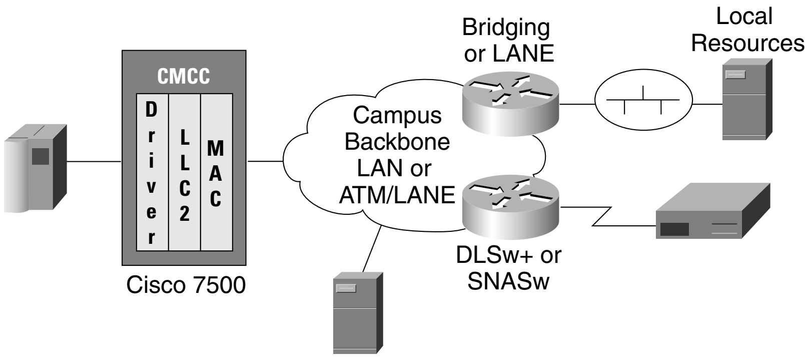

Remote SNA devices can attach to remote Cisco routers and use any of Cisco SNA transport technologies, such as DLSw+, SNASw, Frame Relay Access Support (FRAS), or RSRB to access central site routers. These transport technologies can be running in the CMCC router or in another router that is bridged to the CMCC router. This example is shown in Figure 2-9.

Figure 2-9 Connecting Remote, Router-Attached Resources to the CMCC

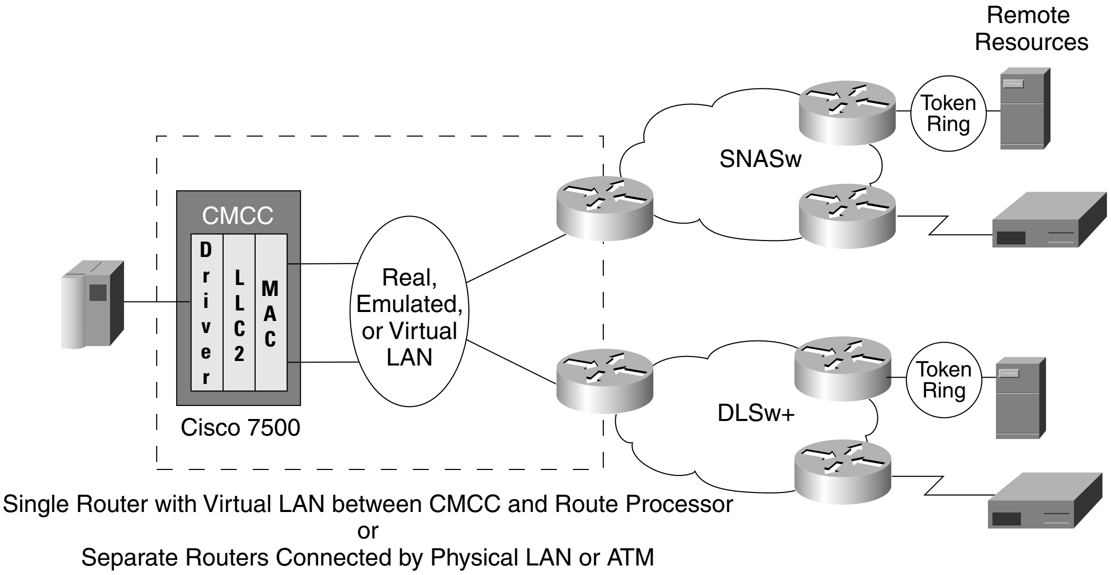
Remote Resources
Token Ring
SNASw
CMCC
Driver
LLC2
MAC
Real, Emulated, or Virtual LAN
Cisco 7500
Token Ring
DLSw+
Single Router with Virtual LAN between CMCC and Route Processor
or
Separate Routers Connected by Physical LAN or ATM

SNA devices that use Qualified Logical Link Control (QLLC) to communicate over X.25 can connect to a Cisco router. Again, you can use either DLSw+ local switching or SNASw to access the CMCC via LLC2. (SNASw or DLSw+ can be running in the CMCC router or in another router that is bridged to the CMCC router.) SNA devices can also communicate directly to a central site Cisco router using RFC 1490 encapsulation of LLC2, as shown in Figure 2-10.

Figure 2-10 Connecting Remote SNA Devices over SDLC, X.25, or Frame Relay

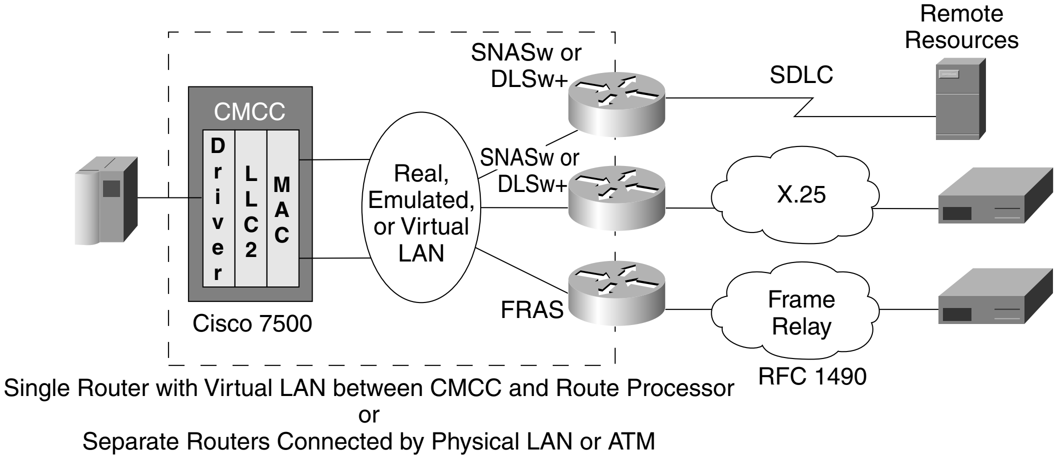

CHAPTER 3

Implementing SNA on the CMCC

This chapter describes the basic requirements to consider before designing and migrating to a CMCC in an SNA network. The information in this chapter will help you answer basic design questions, such as how SNA traffic will travel to the data center and where to place the SNA and WAN functionality to optimize scalability and availability.

Basic Design Considerations

When you design an SNA network using the Cisco CMCC, you must make the following design decisions:

- Determine when to run SNA and WAN functionality on the channel-attached router or on a separate data center router
- Estimate the number of CMCCs required for your network
- Estimate the changes to the mainframe CPU
- Determine if APPN (SNASw) is required
- Determine where to place the DLUR functionality
- Understand and select any appropriate subarea SNA-to-SNASw migration options
- Understand and select any appropriate FEP-to-CMCC router migration options

After you answer the design questions, you can design a network with optimal performance, high availability, and minimal cost.

Accessing Remote SNA Devices

The first design consideration is to determine how SNA traffic will travel to the data center. There are several options:

- Traffic can be bridged
- Traffic can be transported over SDLC, X.25, or Frame Relay
- DLSw+ can transport SNA over an IP backbone
- SNASw can route SNA directly from the branch

These options were described in detail in the previous chapter. You can use these options in any combination. Which of these solutions you choose, and in which combination, depends on the available carrier services, the applications you have or plan to have in your network, and so on. This chapter discusses where to place features to optimize network performance and scalability.

Placement of SNA and WAN Functionality

Central site routers generally provide WAN connectivity, SNA functionality, and mainframe connectivity via the CMCC. You can place all of these functions in a single router. Alternatively, you can have minimal functionality in the CMCC router and place SNA and WAN functionality in other central site routers. The two reasons for not placing functionality in the central site router are scalability and availability in your data center.

Data Center Scalability

Running SNA functionality, such as DLSw+ or SNASw in your channel-attached router, can limit the scalability of your data center solution or increase the cost of the network.

If you run only source-route bridging (SRB) in your channel-attached router and bridge SNA traffic onto one or more CMCCs in the router, the cumulative capabilities of the installed CMCCs are the only limiting factors. The processor in the router uses fast switching for SRB traffic, and the router processor is fast enough to handle traffic coming over many CMCCs. Each CMCC can handle LLC2 processing up to 6000 SNA PUs and can process approximately 5000 packets per second (pps). The capacity of a Cisco 7x00 Series router with multiple CMCCs is additive because the LLC2 processing is contained within the CMCC itself. Because the LLC2 data processing is independent of the router processor, a Cisco 7513 router with 10 CMCCs can process 50,000 pps, which is well within the SRB capability of the router (assuming there are no access lists or filters running).

If your backbone uses transparent bridging, the router uses source-route translational bridging (SR/TLB) to switch packets onto the CMCC. SR/TLB is also fast switched (in Cisco IOS Release 11.2 and later) and the Cisco 7500 Series router can handle more than 25,000 pps. The CMCC can handle a much higher traffic volume than seen in most SNA data centers.

If you place DLSw+ or SNASw/DLUR in the channel-attached router, the number of SNA devices that you can connect is substantially less, up to 4000 SNA PUs if the transaction rate is low and other processor-intensive features are not running in the router. In addition, the transaction rate supported by a route processor is less than the rate supported by the CMCC. In both cases, the limiting factor is the route processor, not the CMCC.

Deploying SNASw and DLSw+ for SNA application routing and WAN transport functionality in separate routers and using the CMCC router for only SRB and IP can allow you to scale your network up to thousands of SNA devices with a single channel-attached router (equipped with one or more CMCCs) and thereby can reduce the total cost of the network.

Determining How Many Channel-Attached Routers and CMCCs Are Required

Function placement plays a role in determining how many channel-attached routers your network requires. Traffic volumes and the number of SNA PUs also play a role. This section provides some guidelines, but you should consult your local systems engineer as well.

There are two limitations you must consider when determining how many channel-attached routers and CMCCs your network requires: the capacity of the CMCC and the route processor capacity.

Selecting CMCC Capacity

Determining how many CMCCs are required depends on the transaction rate, the transaction size, and the number of LLC2 connections. It also depends on the number of hosts to which you will attach the CMCCs and what channel type you have, ESCON or bus and tag. If you are running multiple functions in a single CMCC (such as IP Datagram, TCP Offload, or TN3270 Server), consult your systems engineer for assistance.

The following data assumes that only the CSNA feature is running on the CMCC:[1]

- A CMCC with a single bus and tag can handle up to 6000 LLC2 connections and forward about 5000 pps.
- A CMCC with a single bus and tag can attach to a single host and support up to 32 control units.
- A CIP with two bus and tag daughter cards can attach to two hosts and support 32 control units on each host.
- A CMCC with a single ESCON can attach to between 1 and 32 hosts using either the ESCON Director or ESCON Multiple Image Facility (EMIF).
- A CIP with two ESCON daughter cards can attach to between 1 and 64 hosts by using either the ESCON Director or EMIF.

Other factors can increase the number of CMCCs required, such as availability and redundancy requirements and the number of other features running in the channel-attached router.

Selecting Channel-Attached Router Capacity

If you are running only SRB and IP in the channel-attached router, you can easily place four to five CIPs in a single Cisco 7500 Series router. In this case, availability, redundancy, and risk are the determining factors, rather than performance.

If you run features such as SNASw or DLSw+ in the channel-attached router, the main router CPU, not the CMCC, typically is the limiting factor. A Cisco NPE300 on the 7200 Series router or a 7500 Series router with an RSP4 running DLSw+ can support up to 1100 128-byte data frames per second (at 50 percent CPU utilization). An RSP8 can support at least 2600 data frames per second (at 50 percent router CPU utilization).

Table 3-1 shows that DLSw+ can support data frame rates with TCP encapsulation.

Table 3-1 Processors and Their Approximate Data Frame Rates

Processor	LAN Medium	Frame Data (in data frames per second)
NPE300	Ethernet	1700
NPE300	Token Ring	1900
RSP4	Ethernet	1200
RSP4	Token Ring	1100
RSP8	Ethernet	2100
RSP8	Token Ring	2600

SNASw has similar limitations in the number of PUs. If SNASw or DLSw+ is running in the channel-attached router, a single CMCC can keep up with any SNA traffic the router processor (Cisco 7500 Series router) or main router CPU (Cisco 7200 Series router) can send. In this case, the only reasons to place multiple CMCCs in the router are for redundancy and to handle other functions, such as TN3270 Server or TCP Offload. For medium to large networks, it is more efficient to separate process-switched SNA functionality from the CMCC router.

Other limiting factors can be Open Shortest Path First (OSPF), Frame Relay, or X.25. These features limit the capacity of the router, not the CMCC. However, if these features are running in CMCC-attached routers, they can limit the capacity of the combined solution. This guide does not describe how to determine the router processor limitations of these features.

[1] VTAM may limit the number of control units available for XCA devices.

Attaching the CMCC to a Campus Backbone

FEPs traditionally are attached to the campus network via Token Ring. If you use a channel-attached router, the connectivity choices are much more flexible. Virtually any campus technology can be utilized, including but not limited to Fast Ethernet, Gigabit Ethernet, and ATM.

Data Center Availability

Running multiple functions and CMCCs in a single router can affect availability. If you place all your functionality (CMCCs, DLSw+, APPN/SNASw, Frame Relay traffic shaping, and so on) in a single router, you increase the likelihood of an impact due to a planned or unplanned outage in that router.

For example, if you need to upgrade to a new Cisco IOS Software level so you can use SNASw or the latest DLSw+ features, you must reload the router. (With SNASw, you can design your network to nondisruptively reroute around planned or unplanned outages of a CMCC router.) If you separate SNA functionality from the CMCC router, you minimize the potential for planned or unplanned outages in your CMCC router. SRB functionality rarely requires updating for enhanced functionality. In addition, the fewer functions you run in the CMCC router, the less likely you are to have a failure. However, the tradeoff is that your SNA traffic must now go through an additional hop, creating another potential point of failure.

Some configuration changes, such as changing the maximum transmission unit (MTU) on a router cause the interface cards to restart, including the CMCC. Cisco recommends making changes to the channel-attached routers only during nonproduction hours (realizing that change windows are becoming smaller). Limiting the function running in the channel-attached router minimizes the need for configuration changes.

Regardless of where you place SNA functionality, you can balance your workload across multiple central site devices to minimize the number of sessions disrupted by a single failure. Figure 3-1 compares three functionality placement alternatives.

Figure 3-1 Alternatives for Functionality Placement

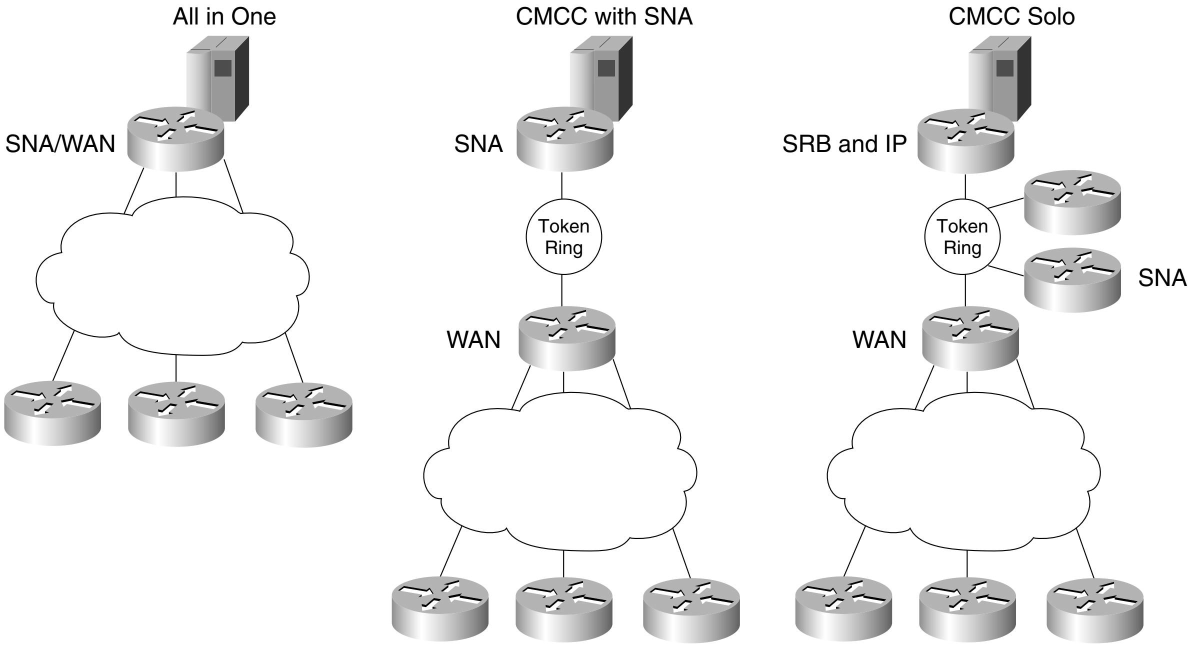

All in One (SNA, CMCC, and WAN)

As shown in Figure 3-1, the first solution is the All in One placement, which requires the fewest central site routers because the CMCC router is also a WAN router with SNA functionality (DLSw+, SNASw, and so on). This solution is reasonable in small networks (30 to 50 branches) that are primarily SNA.

Combined CMCC and SNA

CMCC with SNA, the second solution shown in Figure 3-1, combines a CMCC router with SNA functionality. A separate WAN router is a good solution for small to medium-sized networks (up to 200 remote branches) with a moderate amount of multiprotocol traffic. This solution allows you to segregate multiprotocol broadcast replication from SNA processing.

CMCC

As shown in the CMCC Solo solution in Figure 3-1, bridging to a CMCC router is a good solution for medium to large networks (more than 200 remote branches) that require more than one or two central site routers for SNA. By segregating the SNA and WAN processing from the CMCC-attached router, you can scale the network without buying additional CMCC routers. By minimizing the functionality in the CMCC router, you are maximizing its availability. The SNA functionality can be either in the WAN router or in separate peer routers at the data center. For large networks, using separate SNA and WAN routers enhances scalability and maximizes availability.

Designing for High Availability

High availability is key when accessing SNA applications on a mainframe. This section describes how to achieve high availability by providing alternate data-link paths to access a mainframe and automatic (but disruptive) recovery around failures on a channel gateway. Nondisruptive rerouting around channel gateway failures can be achieved only with SNASw with HPR (available in the Cisco IOS Release 12.0 or later) or TCP (when using the TN3270 Server on the mainframe).

High Availability Using Enterprise Extender

In the case of HPR, nondisruptive rerouting occurs only between the RTP endpoints. Loss of an RTP endpoint is disruptive to the end users' sessions. In most cases, VTAM is one of the endpoints. The other endpoint can be one of the following:

- *In another data center router*—If you place the RTP endpoints in separate (and typically less-expensive) data center routers, you enable nondisruptive rerouting around a channel-attached router failure. In addition, you can balance SNA resources and traffic across a number of data center routers to minimize the impact of a single failure.
- *In the channel-attached router*—Cisco recommends that you avoid placing the RTP endpoint in the channel-attached router because it becomes a single point of failure. The failure of that router is catastrophic (because so many resources use it).
- *At each branch*—The RTP endpoint can be in the branch. Because maintaining large numbers of RTP endpoints places an additional burden on VTAM, Cisco recommends that you measure the impact this placement might have on VTAM. If your network strategy is to move to an IP backbone and isolate SNA to the data center, then the data center router is the preferable location.
- *At the desktop*—Cisco does not recommend extending HPR to the desktop because it increases the workload in the VTAM, thus adversely affecting VTAM performance.

High Availability Using CSNA

Because the CMCC appears as a LAN port and the attaching SNA end system appears as part of a switched major node, the CMCC takes advantage of the redundancy features inherent in LANs and SRB. Cisco DLSw+ provides SRB availability characteristics for devices on Ethernet or SDLC.

In the simplest network, SNA end systems attach to the channel-attached router over a source-route bridged LAN. As you will see, all other networks can be viewed as variations, so this simple network design will be examined first. To follow this section, a basic overview of how SNA works over a SRB LAN is key.

For a LAN-attached end station to gain access to a host over a CMCC, you must configure the end station with the Media Access Control (MAC) address of the CMCC. In addition, you must configure the IDBLK and IDNUM specified in VTAM to match the corresponding value in the end station.

When an end station initiates an SNA connection, it sends an explorer frame (either a TEST or an XID) specifying the MAC address of the CMCC. This explorer is copied by every SRB on the path between the end system and the CMCC. As each bridge copies the frame, it records its bridge number and the next ring number in the Routing Information Field (RIF). If multiple paths exist between the end station and the CMCC, the CMCC will receive multiple copies of the explorer, as shown in Figure 3-2.

Figure 3-2 Explorer Processing on a Source-Route Bridged LAN

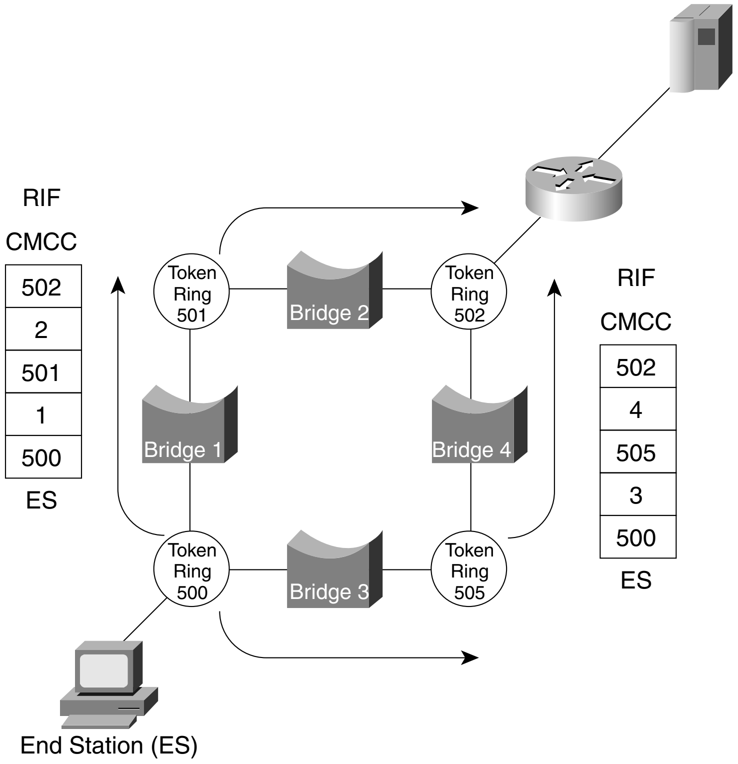

The CMCC responds to each one and sends the response over the same path the explorer took (as specified in the RIF). The end station then selects the route that will be used for the session, typically the one noted in the first explorer response received.

The end station then sends an XID to the CMCC using this SRB path. The CMCC forwards the XID to VTAM. The XID contains the IDBLK and IDNUM and must match an entry defined in VTAM for the session to be established.

Using Duplicate Addresses for High Availability

Source-route bridged LANs support duplicate MAC addresses (as long as they are on different ring segments), because to the end system, they appear as two different paths to one device. The CMCC architecture takes advantage of this characteristic to offer redundancy and load balancing. If a CMCC is out of service for any reason, and another CMCC with the same MAC address is located on a different ring segment, SNA end stations automatically find the alternate CMCC and use it to access the mainframe, as shown in Figure 3-3. In this example, recovery from the loss of a CMCC or channel adapter is automatic but disruptive. Because both CMCCs are in the same router, failure of the channel-attached router is not addressed in this design.

Figure 3-3 Using Duplicate MAC Addresses with CMCCs

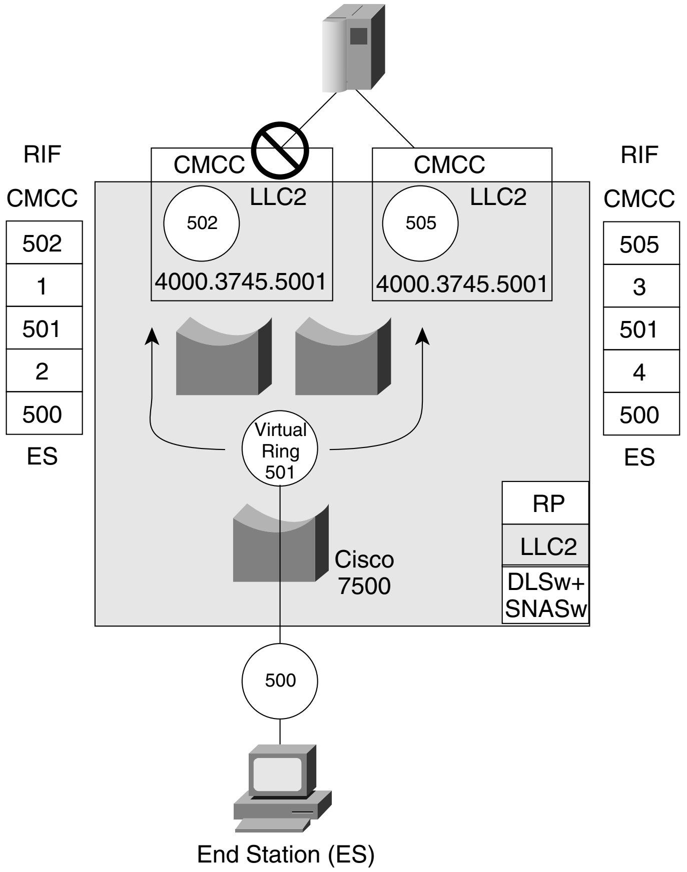

Duplicate addresses are not allowed on a single LAN segment. Duplicate addresses on different segments can be concurrently active. Note that in the case of the CMCC, these segments (as indicated by rings 502 and 505 in Figure 3-3) can be logical ring segments. (In the case of FEPs, these would be physical ring segments.) Other network devices distinguish the duplicate addresses by the RIF used to reach them. The CMCC always uses SRB internally to the channel-attached router, as illustrated by the logical bridges inside the Cisco 7x00 Series router; therefore, multiple CMCCs with the same MAC address can be active as long as they have unique virtual ring numbers. Duplicate CMCCs increase availability by automatically providing redundancy.

In a transparent bridging environment, duplicate addresses are not allowed. However, through the use of DLSw+, Ethernet-attached devices take advantage of the CMCC redundancy described previously. In addition, DLSw+ allows SDLC devices to benefit. In the case of SDLC devices, the MAC address of the CMCC is configured to DLSw+ instead of the SNA device.

You can also use duplicate MAC addresses to load balance traffic across multiple routers equipped with CMCCs and connected to the same VTAM. Figure 3-2 and Figure 3-3 show two possible designs. Both designs provide automatic backup for the loss of a channel-attached router, CMCC, or channel adapter. Load balancing minimizes the number of resources affected by any single outage.

In Figure 3-4, load balancing occurs using standard SRB techniques of the end system. Most end systems select the first path to respond, but eventually the end systems spread over both CMCCs because congestion on a given path through the LAN network, or in a given gateway, slows down the response and leads to the selection of a different path.

Figure 3-4 Load Balancing Using Duplicate MAC Addresses and SRB

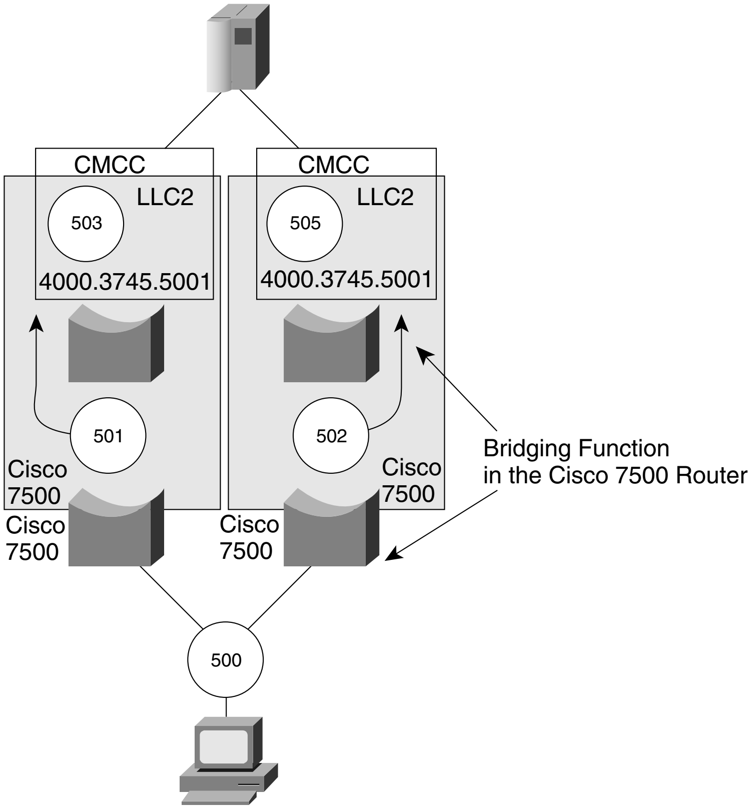

In Figure 3-5, DLSw+ is configured to load balance, which means that each new circuit will alternate, in round-robin fashion, through the list of ports it uses to access the CMCC. If DLSw+ is running in the same router as the CMCCs, DLSw+ views each CMCC as a different port, making load balancing possible.

Note: DLSw+ caches only one RIF per MAC address on a given port; therefore, you cannot load balance across duplicate MAC addresses accessed over a single port, even if they have different RIFs. Cisco recommends using multiple ports.

Figure 3-5 Load Balancing Using Duplicate MAC Addresses and DLSw+

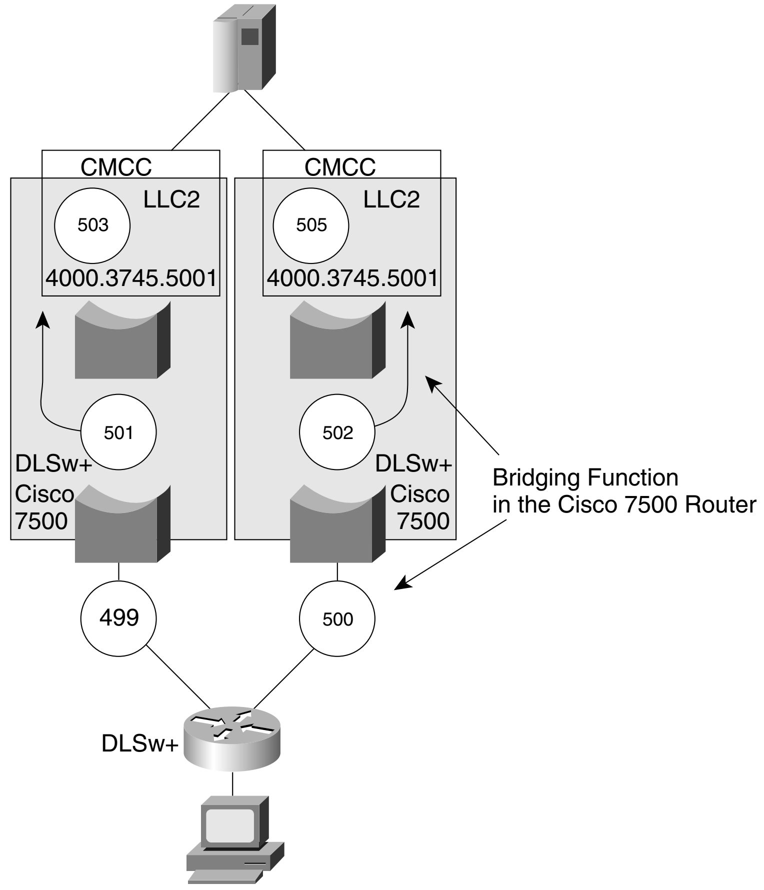

Supporting SNA Channel Protocol with a CMCC

The CSNA software feature offers two commands that you can use to communicate over the channel: **csna** and **cmpc.** When using the **csna** command, the CMCC appears to VTAM as an XCA. When using the **cmpc** command, CMCC communicates to the VTAM using Cisco MultiPath Channel (CMPC). Support for CMPC is available in the Cisco IOS Release 11.3 and later.

The Cisco CMCC running CSNA uses VTAM XCA support. XCA allows one subchannel to support thousands of SNA PUs. XCA is a protocol primarily used for VTAM-to-VTAM, VTAM-to-PU 2, and APPN ISR traffic. HPR is not supported by VTAM using XCA. XCA uses a single half-duplex subchannel to communicate with VTAM.

CMPC is used for VTAM-to-VTAM, VTAM-to-APPN/ISR, and VTAM-to-HPR communication. It requires at least two subchannels for each adjacent SNA PU. CMPC implementation supports one read channel and one write subchannel per adjacent SNA PU. It provides more efficient channel and mainframe utilization than the XCA (or the Channel Data Link Control [CDLC] protocol used by FEPs). However, CMPC can require more configuration than XCA because CMPC supports only one adjacent PU over a pair of subchannels, whereas XCA supports thousands of PUs over a single subchannel. When you implement CMPC, design your network to minimize the number of adjacent SNA PUs and the required definitions.

Supporting SNA Appearance of the Channel-Attached Router

The CMCC does not have an SNA PU appearance. It appears to VTAM as one or more XCA major nodes. (One XCA major node is required for each internal LAN adapter configured with the **csna** command.) A single CMCC can support up to 6000 SNA PUs. Multiple CMCC cards can run in a single router, providing mainframe access for tens of thousands of SNA PUs and LUs. In addition, by using an ESCON Director, the CMCC can connect up to 32 channel-attached mainframes.

Although the CMCC does not have an SNA PU appearance, the router can have an SNA PU appearance. The router appears like an SNA PU 2 when it is configured with a downstream PU (DSPU) concentration or with the service point function, which allows you to manage the router from IBM's Tivoli NetView for OS/390 or Computer Associates' NetworkIT NetMaster.

Migrating from a FEP and Coexisting with a CMCC

Most SNA networks today use FEPs to access their mainframes. A Cisco router solution offers a cost-effective alternative to FEPs in many environments. The question becomes how do you migrate from a FEP to a CMCC? Can a FEP coexist with a CMCC either permanently or during migration?

The easiest and safest way to migrate from a FEP to a CMCC is to use SRB and configure duplicate MAC addresses on both the CMCC and the FEP. This technique requires no changes to the end systems and minimal or no changes to the FEP (assuming the FEP has a Token Ring adapter and is configured with a switched major node). The SRB protocol provides rudimentary load balancing across two mainframe channel gateways and automatic and dynamic backup of one for the other. If the existing FEP does not have a Token Ring card, you can still migrate SDLC devices, one line at a time, by connecting the SDLC line to a router instead of the FEP. The router uses local DLSw+ or SNASw and convert the SDLC to LLC2 for access to a CMCC.

Figure 3-6 shows an example migration from a FEP to a CMCC.

Figure 3-6 Migration from a FEP to a CMCC

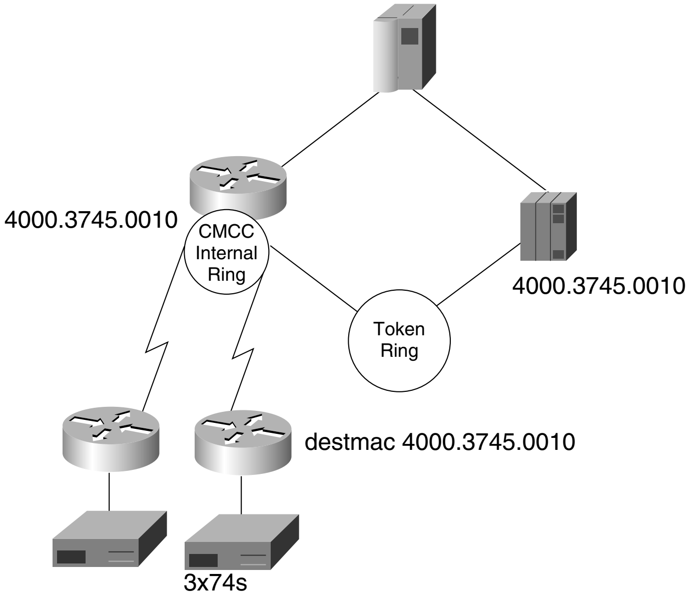

Utilizing Mainframe CPU

A commonly asked question when considering CMCC as an alternative to FEP is what is the impact on mainframe CPU cycles. Replacing a FEP with an XCA channel-attached device (such as a Cisco 7500 Series router with CIP2) has a minimal effect on a given data center capacity plan, and only if the current plan is nearing capacity. Testing has shown that if you replace a FEP with a CMCC, you will see a slight increase (1 to 3 percent) in total mainframe CPU. The increased throughput of the CMCC, however, allows file transfers to occur in less time, freeing the mainframe CPU sooner.

What has confused this issue in the past is the way Resource Monitoring Facility (RMF) measures host CPU usage. RMF does not accurately represent the allocation of CPU cycles to specific tasks. For example, when using the FEP/CDLC path through VTAM, a simple application write operation appears to spend more time in the application address space than in the VTAM address space. Because RMF charges the time spent in VTAM to the application, it makes the VTAM utilization appear very low. When using the XCA path through VTAM, the same application operation spends more time in the VTAM address space and less in the application address space. Because RMF does not charge this time to the application, as it did on the FEP/CDLC case, RMF makes the VTAM utilization appear much higher. What is important from a capacity planning perspective is the difference between the total CPU utilization of the two operations. This difference is what was measured.

In addition, the delta MIPS required to handle a given transaction load is smaller if the transaction rate is higher, partly due to coat-tailing at higher loads, and because the VTAM service request block performs more work under each dispatch (that is, it is able to process more PIUs before releasing the processor).

In one test, a 40-MIPS mainframe (9672-R22) had an average transaction rate of 4500 transactions per minute (75 tps). In this case, replacing the FEP with a Cisco channel-attached router increased host usage by 1.77 percent (0.78 of a MIPS). Running the same host at 70 percent of available cycles and processing transactions at a rate in excess of 11,000 transactions per minute (185 tps) required an extra 3 percent (1.2 MIPS).

Figure 3-7 illustrates the increase in host CPU (delta MIPS) based on the transaction rate in a 9121-982 mainframe. This mainframe is a higher-MIPS machine than the mainframe used in the testing cited in the previous paragraph, so the results will vary slightly.

Note: In addition to a slight increase in mainframe MIPS, migrating from a FEP to an XCA device can increase the memory requirements of the mainframe. You can estimate memory requirements from formulas provided by IBM.

Figure 3-7 Impact on a 9121-982 Mainframe of Migrating from a FEP to a CMCC

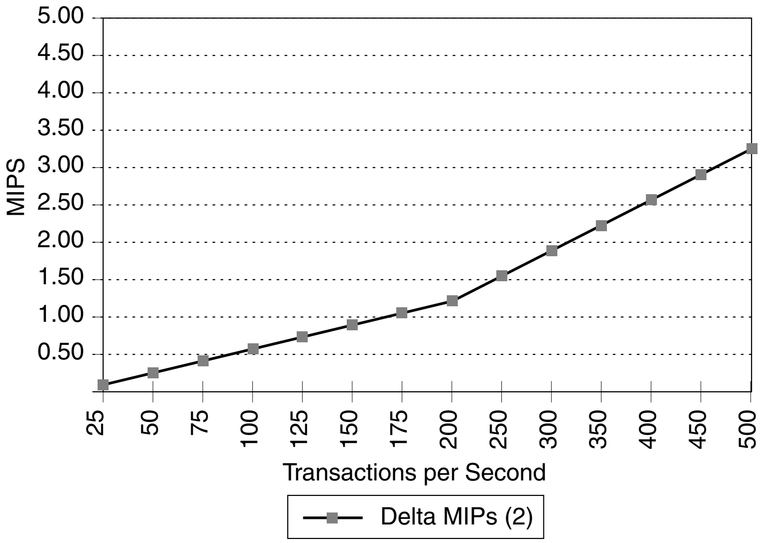

Using SNASw in the Data Center

As enterprises move to an IP infrastructure with SNA and IP applications in the data center, you must replace FEPs to provide adequate TCP/IP support. Channel-attached Cisco CMCC routers provide the connectivity, while SNASw on Cisco routers provides SNA routing between S/390 enterprise servers. If multiple enterprise servers exist, an SNA routing decision must be made. Traditionally, SNA routing decisions were made in the enterprise server VTAM or FEP.

The Cisco IOS Software provides SNA routing with SNASw, so first you must determine whether migrating to a Cisco CMCC solution for SNA requires migrating to SNASw in some portion of your network. (There are many reasons for migrating to SNASw; however, this section only discusses whether SNASw is required to provide SNA routing instead of using FEPs for the same purpose.)

If you currently are using FEPs and running subarea SNA, and you are considering using the CMCC to replace one or more of your FEPs, you may need SNASw in your network. You do *not* need SNASw for SNA routing if any of the following is true:

- You have only one active VTAM image at a time (you may have a second image just for backup).
- Your end users access only a single VTAM; they do not have cross-domain sessions (that is, end users access applications only in a single VTAM LPAR).
- Your end users have sessions with applications in multiple hosts, but SNA sessions use VTAM and Channel to Channel (CTC) for session routing, not your FEPs.
- You have a session manager through which all steady-state session traffic flows.

If you run multiple VTAM images concurrently, and your end users log on to one VTAM and then establish cross-domain application sessions with another VTAM, you are performing SNA routing in the data center. If a FEP is currently performing that SNA routing function, you should consider SNASw in your CMCC network design.

To achieve the benefits of APPN, a minimum number of APPN routers is required although a full NN implementation in the data center router is unnecessary. The goal of an enterprise should be to add sufficient APPN support to provide the needed SNA routing, while minimizing the amount of traffic in the network. SNASw does this with BX, which provides direct routing of data to the correct application host, supports all downstream subarea and APPN devices, and minimizes the scalability and complexity issues associated with a network containing a large number of NNs. SNASw also provides EE, which can be used to more fully integrate APPN into the IP network and data center.

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

CHAPTER 4

Introducing TCP/IP on the CMCC

The role of the mainframe has shifted as businesses increasingly embrace network-centric computing and the architecture of the Web. This computing paradigm is giving mainframes new opportunities. Front-end processes are separated from back-end processes, and content is separated from infrastructure, giving any type of client transparent access to resources on any type of server, as long as both support the Web interface.

According to IDC, more than 87 percent of S/390s were using TCP/IP to access at least some of the data on the mainframe in 2000. Customers are implementing mainframe TCP/IP and evolving toward integrated network infrastructures for end users to access SNA and TCP/IP host applications.

This chapter introduces you to using the CMCC in a mainframe TCP/IP environment. The information in this chapter can help you determine when and where to use the TCP/IP features of the CMCC, as well as the best design for your data center to optimize performance and availability.

This chapter includes the following information:

- The forces driving enterprises to adopt TCP/IP on the mainframe
- An overview of the three TCP/IP features available on the CMCC: IP Datagram, TCP/IP Offload, and the TN3270 Server

Why Customers Require TCP/IP on the Mainframe

Industry prognosticators in the late 1980s were fond of predicting the imminent demise of the mainframe in favor of less-expensive, microprocessor-based alternatives utilizing client/server technologies. The mainframe has shown considerable resilience. Businesses today are processing more mainframe instructions per second (measured in MIPS) than ever before. In short, a huge installed base of applications and an extremely robust, secure operating environment have ensured the survival of the mainframe.

The following requirements are the key reasons why support for TCP/IP on the mainframe is increasing:

- TCP/IP is the basis for all Internet and intranet applications.
- TCP/IP has become the de-facto standard network architecture; interoperability with partners, suppliers, and customers requires a TCP/IP infrastructure.
- The mainframe remains a vital piece in the overall enterprise computing environment.
- The communications software for the mainframe, IBM's CS/390, includes TCP/IP at no additional charge.
- Standards exist to allow TCP/IP-based end systems and SNA devices to access legacy SNA applications over a TCP/IP infrastructure (TN3270 and HPR/IP, respectively).

Virtually every enterprise is faced with the task of implementing new applications that offer Internet and intranet access to customers, suppliers, and partners. TCP/IP, the language of the Internet, is the basis for all these applications. In addition, the enterprise intranet requires a sophisticated set of network appliances and network servers to support security, management, and directory services for these applications. These network-based services are all built using TCP/IP.

As a result of the incredible growth in Internet and intranet applications, TCP/IP has become the standard for interoperability between systems. Electronic data interchange (EDI) used to occur over SNA or other proprietary protocols. Now, Web-oriented protocols like Extended Markup Language (XML) are used to electronically exchange information between applications. Customers, suppliers, and employees all expect to be able to interact with the enterprise using the new, standard user interface—the Web browser.

The mainframe has evolved into an applications and Internet/intranet superserver on an open, standards-based network. The continuing success of the mainframe depends on its ability to reinvent itself as a key platform in the next-generation corporate intranets. This reinvention means attracting new application development and providing active TCP/IP support while maintaining the same level of security and other services afforded by SNA.

IBM has played an active role in ensuring the continued vitality of the mainframe. TCP/IP support, which used to be a separately licensed and priced item, is now included in the mainframe communications subsystem, CS/390. In fact, IBM offers a complete line of TCP/IP applications for the mainframe, including a Web server and an application server. Enterprises automatically have TCP/IP support on the mainframe; they simply have to "turn it on" to leverage the growing set of TCP/IP applications.

Fortunately, enterprises can implement new, TCP/IP-based applications on the mainframe without having to completely rewrite their legacy SNA applications. Standards exist and are supported by a variety of vendors that provide access to SNA applications by both TCP/IP and SNA devices. TN3270 is a standard that permits TCP/IP-based end users to access SNA applications using either standard terminal emulator software or a Web browser. HPR/IP (EE, described in the chapter Introducing SNA on the CMCC) permits SNA devices to access SNA applications over a TCP/IP network.

Enterprises must implement TCP/IP in order to support the applications of the 21st Century. The mainframe, a linchpin of most enterprises, must support TCP/IP in order to maintain its vitality and secure its spot as an Internet/intranet superserver.

Overview of the CMCC IP Datagram Feature

The most common approach to provide CMCC attachment to a TCP/IP mainframe is using the CMCC IP Datagram feature. The CIP and the CPA support both the IP Datagram feature.

Both the host and Cisco router, such as the Cisco 7000, 7200, and 7500 Series, transmit Internet traffic in the form of IP datagrams. When the IP routing software in a host or router needs to transmit a datagram, it consults the routing table on the host or router to determine where to send the datagram. Using the destination IP address, the software looks up the address in the table and selects the next hop to which to send the datagram. The datagram travels from one router to another router until the datagram can be delivered directly across one physical network. The CMCC provides the last hop across the channel to the host.

In IP datagram mode, the mainframe TCP/IP stack performs all the required IP datagram processing. Cisco CMCC routers pass IP datagrams to the CS/390 using the Common Link Access for Workstation (CLAW) or Cisco MultiPath Channel Plus (CMPC+) protocol.

Supporting IP Datagram Using the CLAW Protocol

The CMCC can be configured to use the CLAW protocol for TCP/IP communication across the channel. The CLAW protocol uses two host subchannels (read and write). To minimize host CPU overhead when the system is busy, CLAW uses the full capacity of the channel and provides efficient routing to external adapters. CLAW is a continuously running channel program that polls the CPU for channel program processes, uses two subchannels (one read and one write), and uses 32 logical links to route to a specific application. Logical link 0 is used for the control path, and the other 31 are used for applications. An indicator shows that multiple frames represent one data block. IP packets are encapsulated in CLAW frames with a maximum frame size of 4096 bytes and are transmitted to the channel.

CLAW writes last as long as there is data to send. CLAW reads never end. As mentioned in the previous paragraph, two subchannels are used, one for writing and one for reading. One subchannel is used for all outbound data from the mainframe, while the second subchannel is used for all inbound data to the mainframe.

Supporting IP Datagram Using the Cisco CMPC+ Protocol

You can also configure the CMCC to use the CMPC+ channel protocol for TCP/IP communications across the channel. CMPC+ is the Cisco implementation of IBM's MPC+ feature. The CMPC+ feature (in Cisco IOS Release 12.0(3)T and later) supports the MPC+ features and protocol necessary to support IP. 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. CMPC+:

- Runs on the CMCC
- Supports TCP/IP and HSAS transmission group
- Supports one IP start per CMPC+ group
- Supports one read subchannel and one write subchannel per CMPC+ group. The read subchannel and the write subchannel in an CMPC+ group can be on different physical channels.
- Supports up to 64 KB per I/O block

Up to 64 CMPC+ groups can be configured on a CMCC, depending on memory configuration. CMPC+ can coexist with CMPC, TCP/IP Offload, CLAW, TN3270, and CSNA features.

Overview of the CMCC TCP/IP Offload Feature

TCP/IP Offload was developed as a way to offload CPU-intensive TCP/IP processing from the mainframe to the CMCC to free mainframe CPU cycles for application processing. TCP/IP Offload provides substantial savings in host CPU cycles by reducing mainframe cycles when servicing TCP clients in older version of MVS TCP/IP (before OS/390 Version 2, Release 5) and VM.

Note: The TCP/IP Offload feature is relevant prior to CS/390 Version 2, Release 5.

In the TCP/IP Offload implementation, all TCP/IP processing is performed on the CMCC rather than on the mainframe. The CMCC is channel-attached to the mainframe host. Data for the host arrives at the CMCC. The CMCC processes the IP and TCP packet headers and then passes the data to the host. The offload server passes the data to the API at the host so that the data can be delivered to the TCP/IP application. The only TCP/IP processing performed on the host is that of passing data between TCP/IP and the CLAW interface to the CMCC. The TCP/IP Offload feature uses the CLAW protocol for communicating between the CMCC and the TCP/IP address space.

CMCC Offload with IBM's Transaction Processing Facility

Transaction Processing Facility (TPF) is IBM's high-speed, transaction-oriented operating system that is widely used in airline and hotel reservation systems. Using Cisco CMCC with the TCP/IP Offload feature, TCP/IP users can access today's TPF applications and future applications, leveraging high-speed Web servers on TPF. Today, using Cisco CMCCs with the TCP/IP Offload feature is a high-performance, router-based, channel-attached device that supports TCP/IP applications on a host running TPF.

Overview of the CMCC TN3270 Server

As already described, TN3270 is a standard protocol that allows TCP/IP-based end systems to access mainframe SNA applications. TN3270 is a client/server protocol; the client portion is implemented on the end system and the server is implemented either on the destination system (that is, the mainframe) or in an intermediate gateway device. The TN3270 Server converts the TN3270 traffic into traditional SNA traffic for delivery to VTAM.

TN3270 Server software is provided at no additional charge with IBM's CS/390. However, because the server function entails protocol conversion for each and every data packet headed to or from the mainframe, this operation can quickly consume a large portion of the mainframe CPU cycles if the base of users is large. For this reason, the CMCC offers an optional TN3270 Server. All of the TN3270 Server processing is performed on the CMCC, and the SNA data is passed to the mainframe using the CMCC CSNA feature. For more information on the Cisco TN3270 Server, refer to the *TN3270 Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/tn3270sr/tech/tndg_toc.htm.

CHAPTER 5

Implementing TCP/IP on the CMCC

This chapter describes considerations for designing a TCP/IP network using CMCCs, such as transport alternatives, scalability, and availability. The key issue covered in this guide is where to place the features for optimal network performance and scalability. The choices evaluated in this chapter include IP Datagram (using either CLAW or CMPC+) and TCP/IP Offload. The TN3270 Server design issues and options are discussed in a separate document, *TN3270 Design and Implementation Guide* at www.cisco.com/warp/public/cc/pd/ibsw/tn3270sr/tech/tndg_toc.htm.

Basic Design Considerations

When you are designing a TCP/IP network using the CMCC, you must decide the following:

- The number of channel-attached routers and CMCCs required to support scalability (throughput) requirements
- How to design your TCP/IP network for maximum availability
- How to easily migrate from an IBM 3172 Interconnect Controller to a CMCC

When you answer the design questions, you can design a network with optimal performance, high availability, and minimal cost.

Data Center Scalability

In an SNA environment, the key to determining scalability of the CMCC or other channel controller is the number of active sessions. SNA is characterized by relatively small frames of interactive, screen data.

In contrast, TCP/IP traffic is mixed. Some traffic is interactive traffic with small data frames. However, much of the traffic in a typical TCP/IP environment is from operations such as file transfers that require very high bandwidth and consume vast network resources. Therefore, the primary key to determining scalability of the CMCC in a TCP/IP environment is the overall throughput required. Secondary considerations include function placement and the number of TCP/IP connections.

You can determine the number of CMCCs required based on measuring the amount of traffic and thus the throughput required. This section provides some guidelines; however, you should consult your local systems engineer as well.

Cisco has performed a wide variety of throughput tests in its labs. The results of these tests, which have been confirmed in numerous real customer networks, indicate that the CMCC (running IP Datagram) does not inhibit throughput in any of the configurations tested. (A complete list of the configurations tested is available from Cisco.) In other words, the CMCC is capable of supporting as much IP Datagram traffic as the mainframe is able to generate and the channel connection is able to accommodate.

For example, in one test the CIP was able to sustain 116.8 Mbps throughput.[1] The results, as shown in Table 5-1, indicate a 95 percent utilization of each ESCON channel on the dual-port CIP when using a test application with the limited windowing characteristics observed. The results do not represent the absolute throughput capacity and capability for either the CIP or the RSP4. The CIP CPU utilization is 17 percent, which means that the CIP CPU still has a large amount of bandwidth available for use while transmitting data over two channels.

Table 5-1 shows the actual CIP and RSP4 metrics.[2]

Table 5-1 CIP and RSP4 Metrics

Metric	Utilization
CMCC CPU	17%
ECA0 (Actual Channel Utilization CHPID 1)	95%
ECA1 (Actual Channel Utilization CHPID 2)	95%
RSP4 CPU	1%

Extrapolating from these results, one can conclude that a single Cisco 7513 can support six fully loaded CIP interfaces and five additional LAN/WAN interfaces. The TCP/IP throughput data for a single port varies depending on the channel protocol you use.

Throughput as a Function of Channel Protocols and CMCC Performance

The channel protocol you select affects throughput. In general, using the CMPC+ protocol results in a higher throughput than the use of the CLAW protocol. OS/390 Version 2, Release 8 supports a feature called CLAW packing, which allows more data to be sent in a CLAW frame. IBM recommends using MPC+ in OS/390 Version 2, Release 5 or later.

For TCP/IP throughput, the CIP generally outperforms the CPA. However, the difference is relatively small, as shown in Figure 5-1. The new ECPA4 adapter using the CMPC+ protocol offers the best throughput for a CMCC.

[1] For more information, refer to the white paper “CIP Performance Test: Specific Protocol Throughput Tests” at www.cisco.com/warp/public/cc/pd/ifaa/ifpz/chifpz/tech/cipps_wp.htm.

[2] RMFMON and the Systems Activity Display (SAD) were used to confirm the channel path identifier (CHPID) utilization.

Figure 5-1 TCP/IP Throughput as a Function of the Selected Channel Protocol

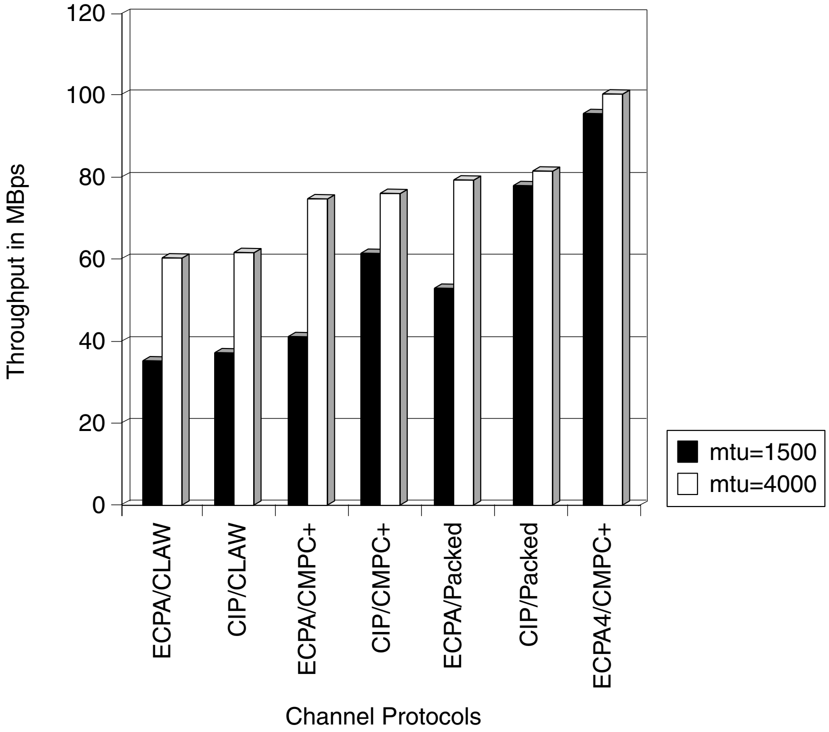

The performance figures in this document are quoted in Mbps. The quotation of throughput figures is not restricted to a particular metric. The use of Mbps (millions of bits per second), or MBps (millions of bytes per second), or KBps (thousands of bytes per second) depends largely on the environment in which the test was performed.

For example, the common terminology used to describe data throughput across IBM channels is MBps. This metric is used partly because the ESCON channel's bandwidth is designated as 10 or 17 MBps. LAN throughput traditionally has been measured in Mbps.

Performance benchmark tests are commonly used to demonstrate the absolute throughput of a given piece of internetworking equipment as described in Request for Comments (RFC) 1944. This document, Benchmarking Methodology for Network Interconnect Devices Status, discusses and defines the tests you can use to describe the performance characteristics of a network-interconnecting device. In addition, RFC 1944 describes specific formats for reporting the test results.

In the tests described in this chapter, the absolute throughput of the internetworking device (the Cisco 7507) has not been tested. The throughput of a single CMCC card and ESCON channels using a single given channel protocol and specific application were tested. (The CMCC can support multiple channel protocols—CSNA, CLAW, CMPC—simultaneously). Furthermore, tests were conducted when sending and receiving frames containing single protocol packets from multiple Token Rings. Consequently, these tests are not considered absolute performance or capability benchmarking tests.

Determining the Number of CMCCs

Determining how many CMCCs are required depends on the transaction rate, the transaction size, and the number of TCP/IP connections. It also depends on the number of hosts you attach to the CMCC and the type of channel you have, either ESCON or bus and tag. If you run multiple functions in a single CMCC (such as IP Datagram, TCP/IP Offload, or TN3270 Server), consult your systems engineer for assistance. The following numbers assume that only the Cisco IP Datagram feature is running on the CMCC:[3]

- A CMCC with a single ESCON can attach to between one and 32 hosts using either the ESCON Director or EMIF.
- A CIP with two daughter cards can attach to between one and 64 hosts by using either the ESCON Director or EMIF.

Other factors that may increase the number of CMCCs required are availability and redundancy requirements and the number of other features running in the channel-attached router.

Figure 5-2 shows CLAW throughput data based on the frame size and fiber length. CLAW has more overhead in the read direction, so Figure 5-2 shows both read and write directions. As the frame size increases, the CLAW channel throughput increases in both read and write directions. The throughput data for CMPC+ follows a similar trend as CLAW. As the frame size increases, so does the throughput.

Figure 5-2 CLAW Throughput as a Function of Frame Size

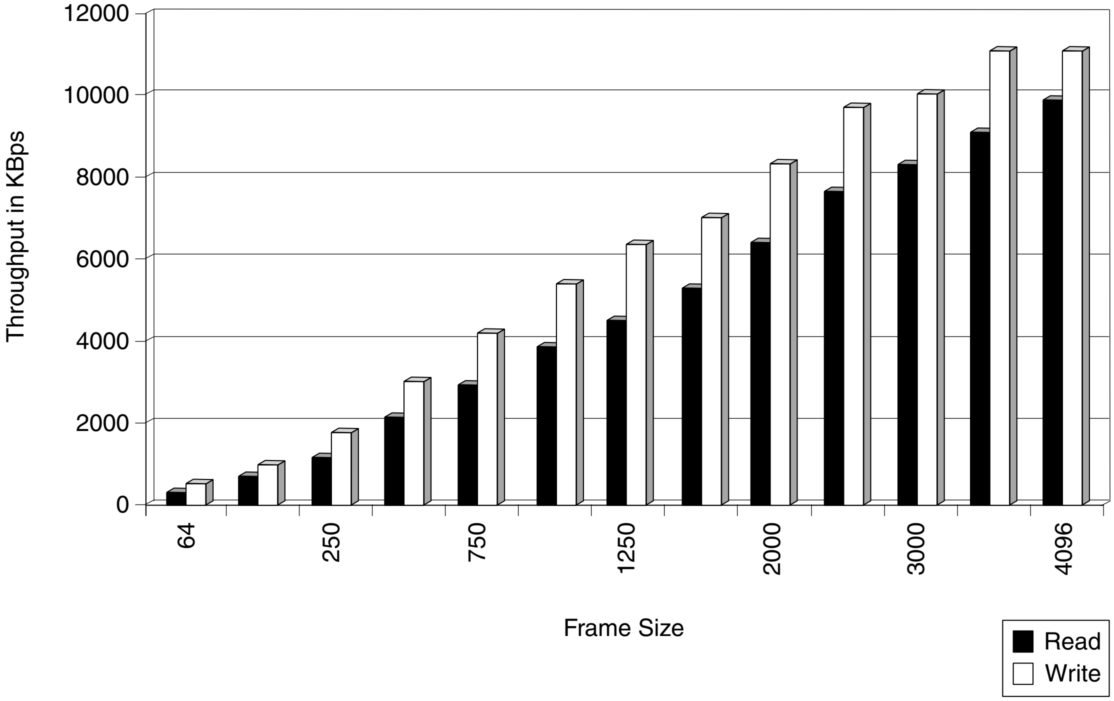

[3] VTAM may limit the number of control units available for XCA devices.

Determining the Number of Channel-Attached Routers

If you run only IP Datagram functions in the channel-attached router, you can place four to five CMCCs in a single Cisco 7500 Series router. Performance is not the determining factor in this case, but rather availability, redundancy, and risk.

Other limiting factors in the router processor can be OSPF, Frame Relay, X.25, or any of the SNA features (for example, DLSw+). These features limit the capacity of the router, not the CMCC; however, if the features run in CMCC-attached routers, they may limit the capacity of the combined solution. This document does not describe how to determine the route processor limitations of these features.

Scalability in a TCP/IP Offload Environment

Before OS/390 Version 2, Release 5, one drawback to using TCP/IP on the mainframe was the amount of mainframe resources required to handle the TCP/IP header, checksumming, and lost packet retransmission. The TCP/ IP Offload CMCC feature alleviates this problem by running the TCP/IP stack on the CMCC. The TCP/IP applications remain on the mainframe.

TCP/IP Offload with a Cisco CMCC significantly reduces the number of cycles needed on the mainframe while still providing high throughput capability. Between 30 and 50 percent of the cycles used by the mainframe stack can be saved by using TCP/IP Offload. Reduced cycle utilization on the mainframe means that the mainframe has more capacity to handle other traffic. The tradeoff for the reduced mainframe utilization is increased utilization on the CMCC. Compared to IP Datagram, the CMCC has a lot more work to do with TCP/IP Offload. Therefore, at some point the CMCC does become a limiting factor for scalability.

TCP/IP Offload can act as an offload host for several ESCON Director-attached mainframes or EMIF LPARs, as well as acting as a single offload host for directly connected ESCON- or bus and tag-connected hosts.

The offload TCP/IP stack can support full ESCON channel throughput of 9 MBps. In addition to offering high throughput, the CMCC, with its 64 MB of memory, maintains 7500 Telnet sessions with a 2-KB window or 500 FTP sessions with a 32-KB window. It simultaneously communicates to 32 host connections through the same ESCON adapter.

Data Center Availability

This section describes the channel-attached routers and CMCCs required to design your data center for maximum availability. Availability is the degree to which your data center resources are available to process transactions. This section provides some guidelines; however, you should consult your local systems engineer as well. Choices include using virtual IP addressing for redundancy and using MultiNode Load Balancing (MNLB). Another important availability issue is the prevention of the loss of traffic. Using Cisco IOS quality of service (QoS) features, the potential for the loss of mainframe traffic can be minimized.

Using the Virtual IP Address Feature for TCP/IP Redundancy

Virtual IP Address (VIPA) is a feature of OS/390 that allows you to configure a single virtual IP address on the mainframe. The network routing protocols Routing Information Protocol (RIP), called ROUTED on the mainframe for an IBM TCP/IP stack, and OSPF dynamically locate an alternate route to the VIPA (configured on the host TCP/IP stack definitions) through routing updates, as shown in Figure 5-3.

Figure 5-3 Redundancy Using the Virtual IP Address

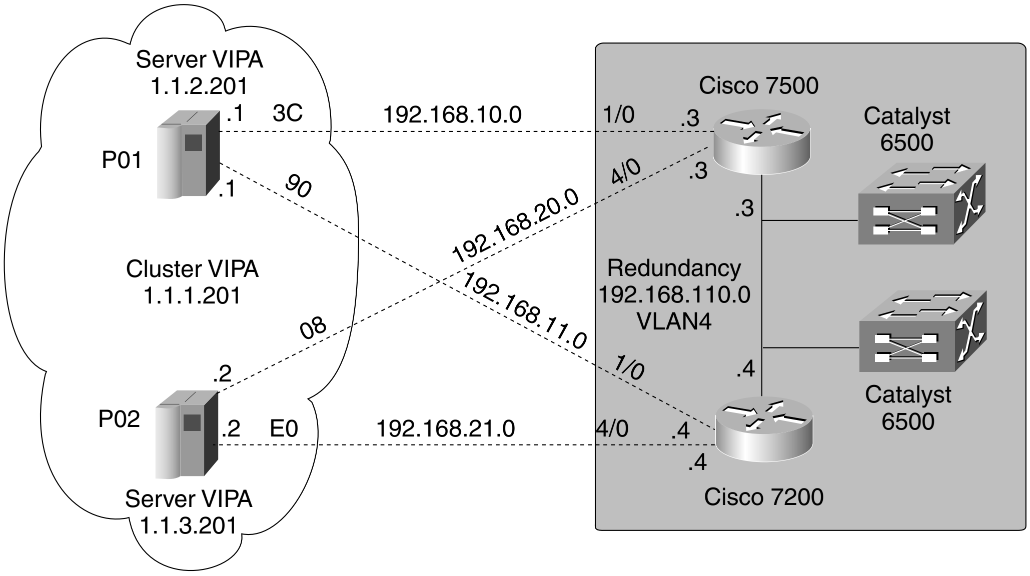

Without the VIPA feature, if the CMCC adapter fails, alternate route information is often ignored because as long as the router's interface on the subnet is active, it sends the packet directly over the subnet.

VIPA solves this problem by creating a virtual subnet within the mainframe TCP/IP address space. When a remote host sends a packet to the virtual IP address, it is forced to use IP routing regardless of whether the host is on a directly attached subnet or not. The remote host must use its routing tables to decide on the best path to the virtual IP address.

Implementing Cisco TCP/IP Sysplex Solution: MNLB

With the rapid migration from SNA- to IP-based networks, Web access to mission-critical mainframe applications residing on host server platforms is essential. The host server cluster must scale to meet future growth requirements and be easy to maintain and support without disrupting application availability to the end user. In addition, the host server components should provide feedback on the network to be used as the basis of load-balancing decisions. MNLB is a Cisco IOS Software feature that provides easy access to IBM TCP/IP application servers residing on multiple hosts. With this feature, you do not need to keep track of which TCP/IP stack is running a particular application or to which host it is best to connect.

MNLB is designed for loosely coupling individual computing servers that balance the workload across the systems. This balancing is transparent to the requesting clients. The architecture employs an IP-based feedback mechanism enabling continuous adjustment of load-balancing decisions. The Cisco Appliance Services architecture (CASA) customizes routing in neighboring IP routing engines (or forwarding agents) under the direction of an MNLB Services Manager.

The MNLB architecture does not require all inbound traffic for a server cluster to pass through a single load-balancing engine. It enables a combination of fast forwarding agents (routers or IP-capable switches) and load-balancing and backup managers to synchronize and control traffic.

The MNLB design allows multiple MNLB routers to distribute load across multiple hosts in a cluster. A global virtual IP address is assigned to the cluster, which must be configured to all hosts as well as the MNLB routers. It allows the hosts to recognize the packets addressed to them and allows the MNLB routers to identify the packets that must be intercepted and rerouted.

The MNLB Services Manager provides a synchronization point for the assignment of affinities and maintenance of the affinity cache for clients and servers. An optional forward IP address identifies the new destination for the packet if it is not sent to the manager. The Services Manager assigns affinities using criteria specific to the network function. For example, the Services Manager uses load management information either sent by the host or statically configured to balance workloads.

A load-sensing agent interacts with the IBM Workload Manager to obtain host load information. This interaction allows the MNLB Services Manager to connect to each host in the cluster and retrieve load metrics that are used in balancing affinities across the cluster. Beginning with zOS Version 1, Release 2, the IBM Sysplex Distributor performs the functions of the MNLB Services Manager. For earlier releases, the MNLB Services Manager function is performed by the Cisco LocalDirector.

Forwarding agents can intercept packets that match local cache affinity and process them as instructed. If a matching affinity is not found, the packet is compared against the wildcard affinities to find managers that are interested in this type of packet. If no appropriate wildcard affinity is found, normal IP routing prevails. Generally, a manager uses the wildcard affinity to be informed of flows. When a manager has determined how a flow should be handled, it sets a full cached affinity so that subsequent packets for that same flow can be offloaded to the forwarding agent.

In the case of load balancing, full affinity is used to identify the server that is to receive the data. However, a wildcard affinity is used to define packet criteria and to identify the manager that makes the balancing decision for the IP packets that match wildcard criteria.

Using the Cisco IOS QoS Features

Cisco enables system administrators to prioritize mainframe traffic over other IP traffic. When you use IP precedence with a Cisco infrastructure, it improves the end-to-end network response time by leveraging Cisco IOS QoS[4] features, such as Weighted Fair Queuing (WFQ) and Weighted Random Early Detection (WRED).

These features ensure that mainframe IP traffic takes precedence over other traffic and minimizes the chance of mainframe traffic being dropped during periods of congestion. This is especially important for SNA sessions that are being supported by TN3270 or HPR/IP. The SNA protocol is designed so that sessions can be dropped if an expected response is not received within a given amount of time. By judiciously utilizing QoS features, you can minimize the possibility of lost SNA sessions due to dropped traffic.

Migrating from an IBM 3172 Interconnect Controller

The IBM 3172 Interconnect Controller, introduced in the 1980s as a specialized PC server for connecting TCP/IP LANs to the mainframe, enjoyed considerable success in the late 1980s and early 1990s. IBM added SNA support to the original TCP/IP support, and it became a common LAN-to-mainframe connectivity device. As a result, many enterprises continue to have an installed base of IBM 3172 Interconnect Controllers in the data center.

[4] For more information about QoS, refer to the white paper "Delivering Predictable Host Integration Services" at www.cisco.com/warp/public/cc/so/neso/ibso/ibm/s390/phost_wp.htm.

However, the IBM 3172 is based on dated technology and can no longer provide the throughput and availability characteristics required by the modern enterprise data center. The CMCC is a complete functional replacement for the IBM 3172 that provides vastly improved scalability and availability. Because of this, enterprises with multiple IBM 3172s can simplify their environments by migrating to a CMCC solution.

The CMCC implements the CLAW protocol for both IP Datagram and TCP/IP Offload features. This is the protocol used for the TCP/IP Offload feature offered on the IBM 3172. Therefore, enterprises that have implemented this feature on the IBM 3172 can very easily migrate to the CMCC by simply ensuring that the CMCC definitions match those of the existing interconnect controllers. Host definitions for IBM 3172s that are operating in an equivalent manner to the IP Datagram support on the CMCC need to change, but the magnitude of the change is very small.

CHAPTER 6

Migration Scenarios

This chapter shows the basic configuration of several of the most common networks. The chapter covers network design and explains why and when to use a particular network design. It briefly describes how to migrate from, or coexist with, a FEP for each of the sample networks. In some cases, before and after pictures of the network and step-by-step configuration instructions are included.

This chapter includes the following scenarios:

- Scenario 1—Replacing a FEP with a single CMCC on a single host
- Scenario 2—Replacing a FEP with a redundant CMCC on a single host
- Scenario 3—Replacing a FEP with a single CMCC on multiple hosts
- Scenario 4—Combining SNASw with DLSw+
- Scenario 5—Migrating to SNASw only
- Scenario 6—Migrating to TCP/IP across CLAW
- Scenario 7—Migrating to TCP/IP across CMPC+

To use the scenarios, you must include the VTAM definitions and configure your routers, which are discussed in the following sections.

Using SNA Communication over CSNA

SNA nodes communicate with the CMCC using LLC2, a connection-oriented data-link protocol for LANs. An LLC2 stack on the CMCC card communicates with either the adjacent SNA device (over a physical Token Ring) or to DLSw+ or SNASw running in the channel-attached router, as illustrated in Figure 6-1.

Figure 6-1 Communication between CSNA in the CMCC and SNA Nodes

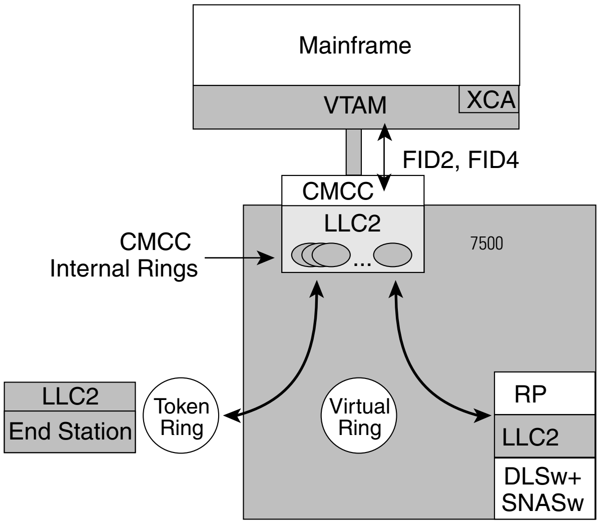

The CMCC running CSNA can support multiple internal LAN interfaces, each appearing as a LAN port to the VTAM. Although VTAM supports a maximum of 18 LAN ports, only a single LAN port is required. CSNA also supports up to 256 open LLC2 service access points (SAPs) per LAN port.

Using VTAM Definitions

The CMCC running CSNA is not an SNA-addressable node, because it has no PU or LU appearance. CSNA is defined to the host control program (MVS or VM) as a channel-to-channel machine (an IBM 3088). CSNA provides VTAM with a physical connection to the LAN through a subchannel.

To enable VTAM communication over the CMCC to SNA devices, you must configure an XCA major node and a switched major node to VTAM. The XCA major node allows VTAM to communicate with the CMCC, and the switched major node definition allows SNA devices to communicate with VTAM over the CMCC.

XCA Major Node Definition

Define an XCA major node for each connection (port) between the VTAM and a CSNA. A single XCA major node can support up to 4096 LLC2 connections, although better results are achieved with 3000 or fewer LLC2 connections per XCA major node. If more LLC2 connections are needed, define additional XCA major nodes as well. You can configure multiple XCA major nodes for availability, with each node pointing to a different CMCC.

The CSNA feature is defined to the host control program (MVS or VM) as being a channel-to-channel adapter (CTCA) or machine; for example, an IBM 3088. VTAM identifies the CSNA gateway through a combination of the following:

- ADAPNO—Adapter number
- CUADD—Subchannel address
- SAPADDR—SAP address

The following configuration provides an example:

```
XCANAME VBUILD TYPE=XCA  ** EXTERNAL COMMUNICATION ADAPT**
PORTNAME PORT ADAPNO=?,  ** RELATIVE ADAPTER NUMBER ** X
            CUADDR=???,  ** CHANNEL UNIT ADDRESS ** X
           MEDIUM=RING,  ** LAN TYPE ** X
              SAPADDR=4  ** SERVICE ACCESS POINT ADDRESS **
GRPNAME GROUP ANSWER=ON, ** PU DIAL INTO VTAM CAPABILITY ** X
        AUTOGEN=(5,L,P), ** AUTO GENERATE LINES AND PUS ** X
             CALL=INOUT, ** IN/OUT CALLING CAPABILITY ** X
                DIAL=YES,** SWITCHED CONNECTION ** X
          ISTATUS=ACTIVE ** INITIAL ACTIVATION STATUS **
```

Switched Major Node Definition

Configure one or more switched major nodes. Within a switched major node definition, configure every SNA PU that will access VTAM through the CMCC. For each PU, configure its associated LUs. Many networks today already include SNA devices defined in a switched major node. For example, if the devices attach to a FEP over Token Ring, the devices are defined as part of a switched major node. In this case, the only change is to add the XCA major node.

The following configuration provides an example:

```
SWMSNAME VBUILD       TYPE=SWNET,       **      X
                      MAXGRP=14,        **      X
                      MAXNO=64          **
PUNAME PU             ADDR=01,          **      X
                      PUTYPE=2          **      X
                      IDBLK=???         **      X
                      IDNUM=???         **      X
                      ISTATUS=ACTIVE    **      X
LUNAME1  LU           LOCADDR=02
LUNAME2  LU           LOCADDR=03
LUNAME3  LU           LOCADDR=04
LUNAME4  LU           LOCADDR=05
LUNAME5  LU           LOCADDR=06
```

Configuring Routers

You must configure the router to:

- Bridge the traffic from a physical LAN or a router component (DLSw+, SRB, SR/TLB, and so on) onto the router virtual ring
- Bridge the data from the router virtual ring to one of the CMCC internal rings, or connect a data-link user (APPN or DSPU) to one of the CMCC internal rings
- Connect the CMCC to VTAM

Figure 6-2 shows the major configuration parameters of CMCC and Token Ring interfaces and how they are logically combined using the source-bridge definition. The CMCC ring is referred to as an internal ring. The Route Switch Processor (RSP) ring is referred to as a virtual ring.

Figure 6-2   Using Virtual Rings to Provide Connectivity

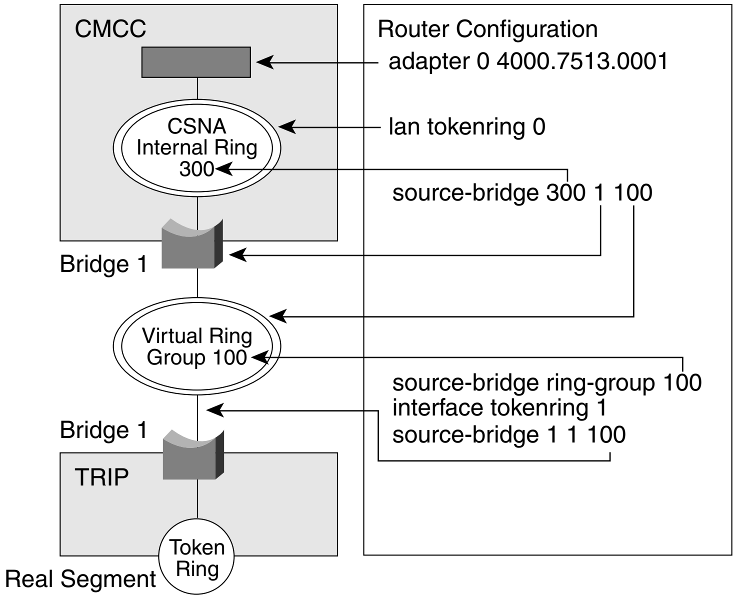

Configure an adapter on the CMCC to associate with the XCA major node definition. For each adapter you configure, CSNA creates an internal Token Ring. A virtual bridge connects the CSNA internal ring to a virtual ring group in the router. The Token Ring Interface Processor (TRIP) is also configured to connect to the same virtual ring group as the CMCC.

Understanding Configuration Relationships in the ESCON Environment

Figure 6-3 shows the relationship among router configuration, VTAM parameters, and MVS IOCP generation commands when the CMCC connects via an ESCON Director.

Figure 6-3   Configuration Relationship in an ESCON Environment

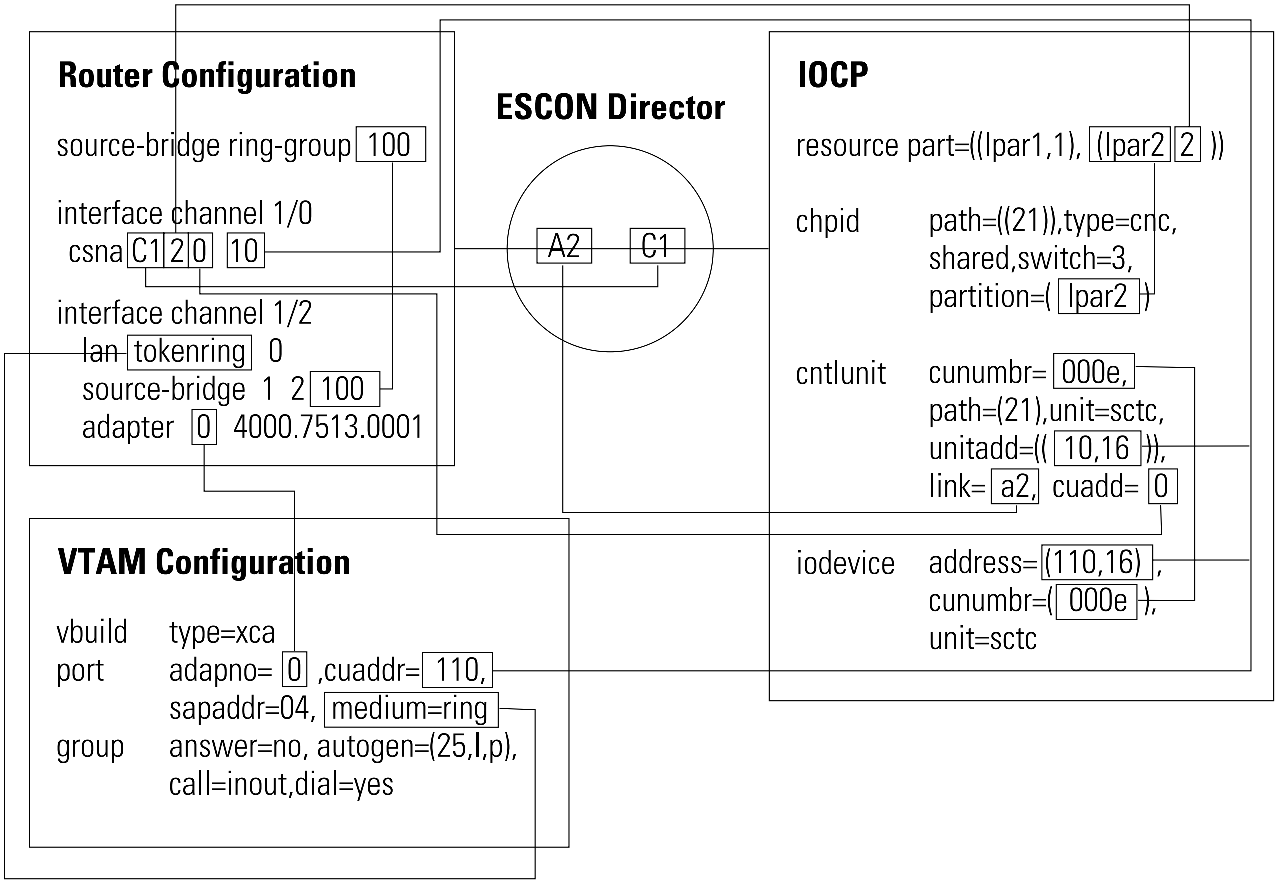

```
source-bridge ring-group 100

interface channel 1/0
 csna C1 2 0  10

interface channel 1/2
   lan tokenring  0
   source-bridge  1  2 100
   adapter  0  4000.7513.0001
```

```
resource part=((lpar1,1), (lpar2 2 ))

chpid       path=((21)),type=cnc,
            shared,switch=3,
            partition=( lpar2 )

cntlunit    cunumbr= 000e,
            path=(21),unit=sctc,
            unitadd=(( 10,16 )),
            link= a2,  cuadd= 0

iodevice    address=(110,16) ,
            cunumbr=( 000e ),
            unit=sctc
```

```
vbuild     type=xca
port       adapno= 0 ,cuaddr= 110,
           sapaddr=04, medium=ring
group      answer=no, autogen=(25,l,p),
           call=inout,dial=yes
```

Understanding Configuration Relationships in the Bus and Tag Environment

Figure 6-4 shows the relationship among router configuration, VTAM parameters, and MVS IOCP generation commands when the CMCC connects via bus and tag.

Figure 6-4   Configuration Relationship in a Bus and Tag Environment

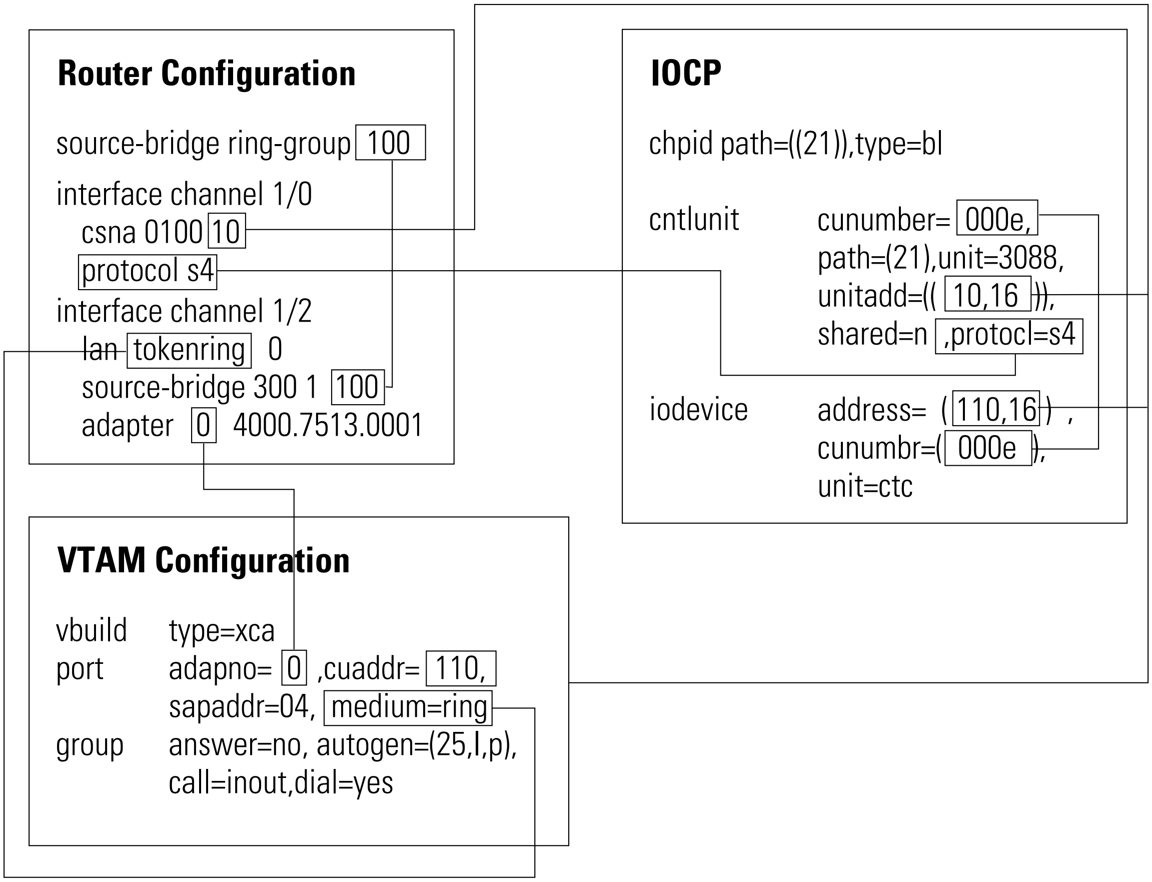

Router Configuration

```
source-bridge ring-group 100
interface channel 1/0
  csna 0100 10
  protocol s4
interface channel 1/2
  lan tokenring  0
  source-bridge 300 1 100
  adapter  0  4000.7513.0001
```

IOCP

```
chpid path=((21)),type=bl

cntlunit      cunumber= 000e,
              path=(21),unit=3088,
              unitadd=(( 10,16 )),
              shared=n ,protocl=s4

iodevice      address= (110,16) ,
              cunumbr=( 000e ),
              unit=ctc
```

VTAM Configuration

```
vbuild    type=xca
port      adapno= 0 ,cuaddr= 110,
          sapaddr=04, medium=ring
group     answer=no, autogen=(25,l,p),
          call=inout,dial=yes
```

Scenario 1—Replacing a FEP with a Single CMCC on a Single Host

The first scenario describes a network that replaces a FEP with a CMCC. As shown in Figure 6-5, a single mainframe exists in this network. Historically, IBM SNA networks were built using the IBM FEP, and remote terminals were connected via SDLC links. In the Before scenario, a second FEP was in place only for backup. In the After scenario, one FEP is replaced with a channel-attached router with a CMCC. Both the CMCC and the remaining FEP have the same MAC address. Eventually, the second FEP also will be replaced, but for now it provides SNI connectivity to a supplier and functions as a backup to the CMCC. DLSw+ is used to transport SNA traffic from remote sites to the central site. When data reaches the headquarters site, DLSw+ sends the traffic to the CMCC, which is the first to respond to explorers. In the event the CMCC is not available, the FEP is used automatically.

Figure 6-5 Single CMCC to Single Host

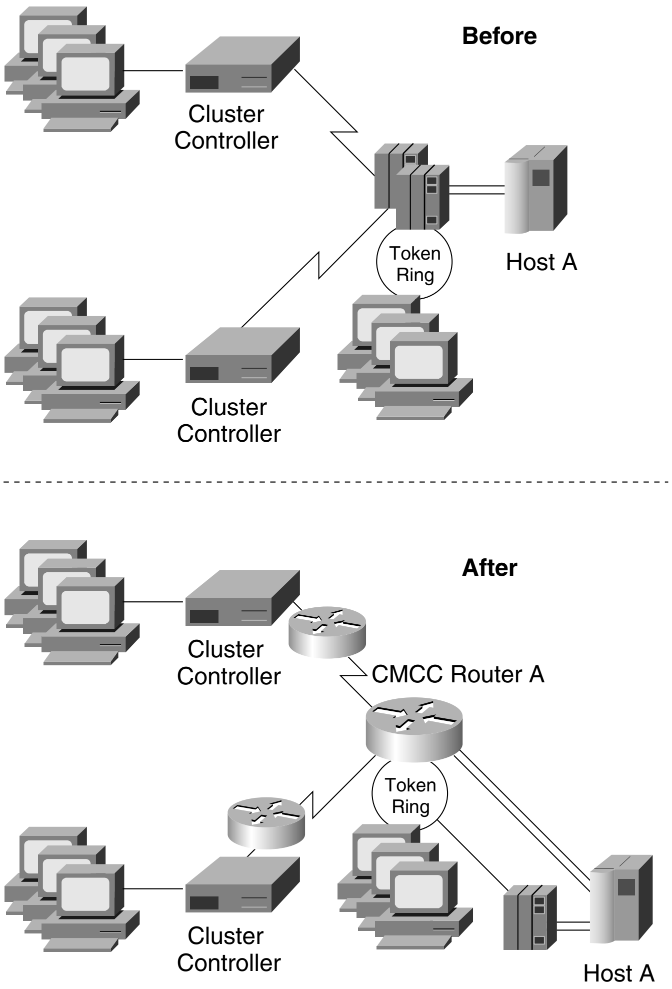

Reasons for Change

The FEP was at capacity and the company preferred to use its IS dollars on technology that would carry the company into the future while addressing today's requirement. In addition, the Cisco channel-attached router replacing the leased FEP would pay for itself in 18 months—with savings coming from lease costs and monthly NCP licensing costs. Migrating from an SDLC/FEP network to a LAN/channel-attached router network simplified SNA system configuration significantly and reduced the downtime for planned outages. Finally, this infrastructure enabled the customer to use TCP mainframe applications in the near future.

Design Choices

This customer opted to combine SNA functionality (DLSw+) and WAN connections in the CMCC router because the network was very small (25 sites). The design provided a very safe fallback to the FEP, but at the same time enabled SRB dynamics and configuration simplicity.

XCA Major Node Configuration

```
XCANODE    VBUILD     TYPE=XCA
PRTNODE    PORT       ADAPNO=0,CUADDR=770,SAPADDR=04,MEDIUM=RING,TIMER=30
*
GRPNODE    GROUP      ANSWER=ON,            X
                      AUTOGEN=(100,L,P),    X
                      CALL=INOUT,           X
                      DIAL=YES,             X
                      ISTATUS=ACTIVE
```

Router Configuration

```
!
source-bridge ring-group 100
!
interface tokenring 1/0
-   no ip address
-   no ip route-cache
-   ring-speed 16
-   source-bridge 200 1 100
!
interface Channel1/0
-   no ip address
-   csna 0100 70
!
interface Channel1/2
-   no ip address
-   no keepalive
-   lan TokenRing 0
-   source-bridge 300 1 100
-   adapter 0 4000.7000.0001
!
end
```

Implementation Overview

The first step is to implement DLSw+ from the remote site to the central site and to change the FEP access from SDLC to Token Ring. As part of this step, configure the VTAM switched major nodes. Next, perform the following steps to enable the CMCC in this configuration:

Step 1. Perform IOCP generations to configure the channel definitions, as shown in Figure 6-5.

Step 2. Configure the VTAM XCA major node.

Step 3. Configure the attached router with the CMCC definitions and bridge traffic from the internal ring group to the CMCC virtual ring.

Step 4. Vary the channel online (**Vary E00,ONLINE**).

Step 5. Confirm the CMCC is online (**Display U,,,E00,1**).

Step 6. Activate the VTAM XCA (**Vary NET,ACT,ID=name_of_member**).

Scenario 2—Replacing a FEP with a Redundant CMCC on a Single Host

Initially this site had an IBM 3745-410 running in twin-standby mode to provide better network resiliency. In this case there is one active NCP while the second one is in standby mode. The second NCP takes over only if the first NCP has problems. This allows quick recovery from storage-related failures and from a CCU hardware check. Note that the idle CCU is inactive unless a failure is detected. With the inclusion of duplicate Token Ring addressing, this design provides another level of network redundancy.

Optionally, the IBM 3745-410 could be configured in twin-backup mode, where each CCU controls approximately half the network. It is the equivalent of having two IBM 210s running at half capacity. If there is a failure in one CCU, the other takes over, just as in the first example. However, only half the resources are impacted, resulting in a faster recovery.

Regardless of the current configuration, the use of CSNA on two Cisco 7500 Series routers with one or more CMCC cards can provide better load sharing and redundancy features, as described in the Designing for High Availability section.

The After scenario is designed without a single point of failure in the network. The redundant CMCC to a single host scenario is often used when the end systems cannot afford the downtime of a failure. For many companies that require online access to provide 24-by-7 customer support, the loss of host access for even a short period can incur a significant loss in both income and credibility. It is important for these networks to implement a solution that avoids or minimizes the amount of downtime due to network problems. Also, for these companies the redundancy option provides the necessary network configuration to perform maintenance or configuration changes with minimal impact on the end-system users.

Providing redundancy to the single CMCC to single host solution is quite straightforward. In Figure 6-6, two Cisco 7500 Series routers, each with a CMCC, are deployed in place of the IBM 3745-410. In this example both CMCCs have the same virtual MAC address. When one router is unavailable, the SNA end system automatically finds the backup router using standard SRB protocols. Note that in both the Before and the After networks, the loss of a channel-attached gateway is disruptive.

Figure 6-6 Redundant CMCCs to Single Host

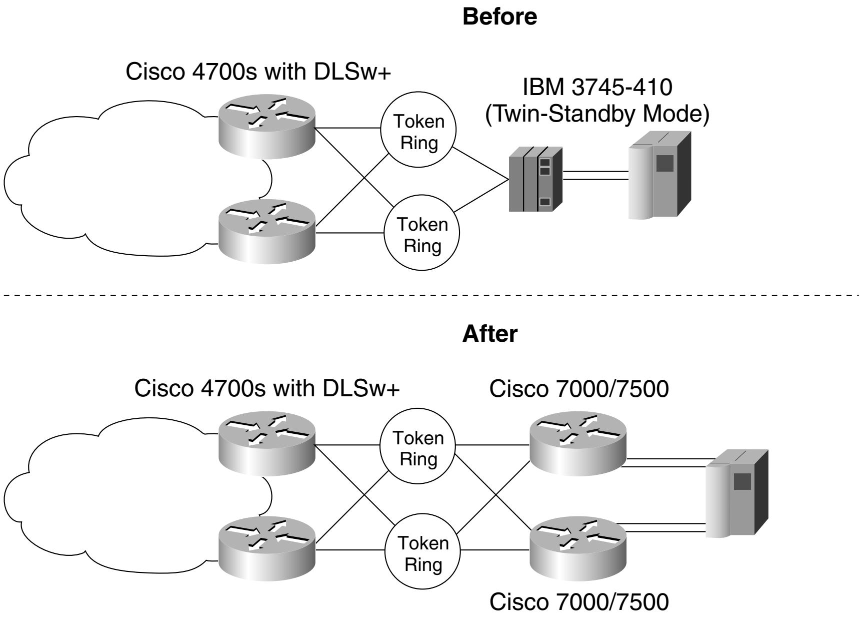

Reasons for Change

The IBM 3745-410 did not have the capacity to support the entire network if one of the processors was down. During outages, the entire network slowed down. To address this problem with more FEPs was not cost-effective. In addition, this enterprise was considering migrating to FDDI, which the IBM 3745 does not support. With the Cisco channel-attached routers, the company could migrate its campus to FDDI, ATM, or Gigabit Ethernet in the future.

Design Choices

In this network they opted to separate DLSw+ from the channel-attached router, thus minimizing both scheduled and unscheduled outages in their network. Also, they already had DLSw+ installed in these routers before they installed the CMCCs, which simplified migration. Finally, as their DLSw+ routers (Cisco 3600s) reach capacity, it would be less costly to add a Cisco 3600 Series router than a Cisco 7500 Series router with a CMCC. Either of the channel-attached routers could handle their entire capacity today, and if the network were to grow, they would have sufficient slots in their Cisco 7500 Series routers to add CMCCs.

The network uses load balancing across central site DLSw+ routers and duplicate Token Rings to ensure there is no single point of failure, as shown in Figure 6-7.

Figure 6-7 Dual Routers with Duplicate MACs

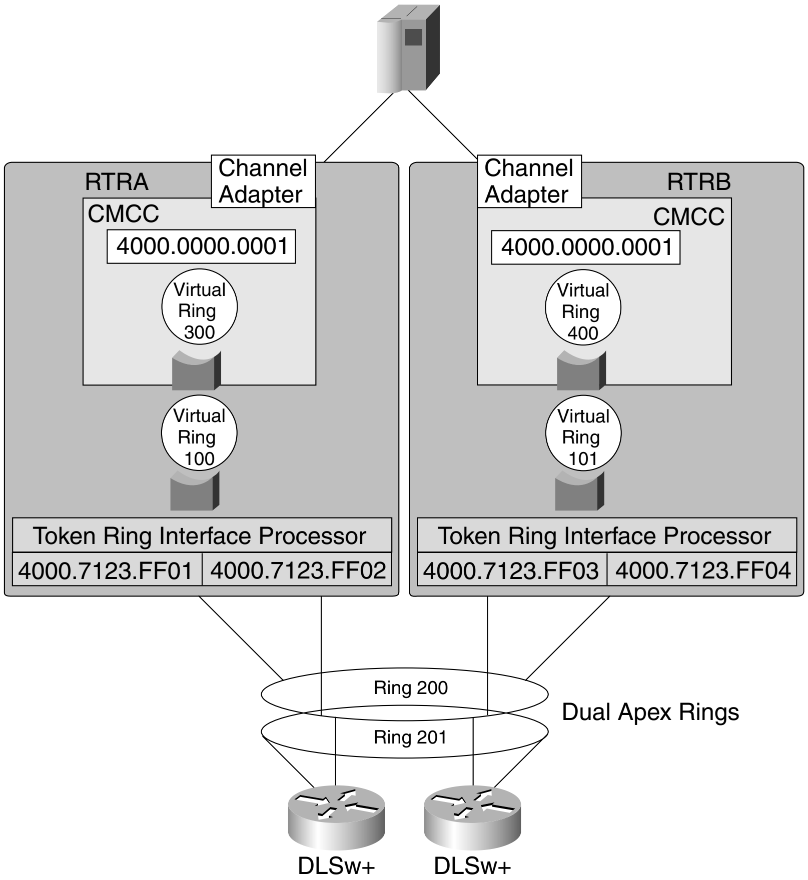

Router Configuration

This configuration uses the same MAC address on internal Token Ring LANs of two different routers:

```
RTRA
!
source-bridge ring-group 100
int tok 0/0
source-bridge 200 1 100
int tok 0/1
source-bridge 201 2 100
!
interface Channel1/0
-   no ip address
-   csna 0100 70
!
lan TokenRing 0
-   source-bridge 300 1 100
-   adapter 0 4000.0000.0001
!

RTRB
!
source-bridge ring-group 101
int tok 0/0
source-bridge 200 1 101
int tok 0/1
source-bridge 201 2 101
!
interface Channel1/0
-   no ip address
-   csna 0100 80
!
lan TokenRing 0
-   source-bridge 400 1 101
-   adapter 0 4000.0000.0001
!
```

Scenario 3—Replacing a FEP with a Single CMCC on Multiple Hosts

This scenario reflects a legacy SNA network with several remote sites connected via SDLC links to cluster controllers. Also, a high-speed line was connected to a remote IBM 3745 at a site that demanded high-speed connection back to the mainframe, but had more remote users than a cluster controller could support. This enterprise also had a separate multiprotocol network running in parallel.

At the data center, there are two VTAM's. One is used primarily for production and the other for testing. There is little, if any, cross-domain routing. Figure 6-8 shows the Before and After networks.

Figure 6-8 Replacing a Single FEP with a Channel-Attached Router

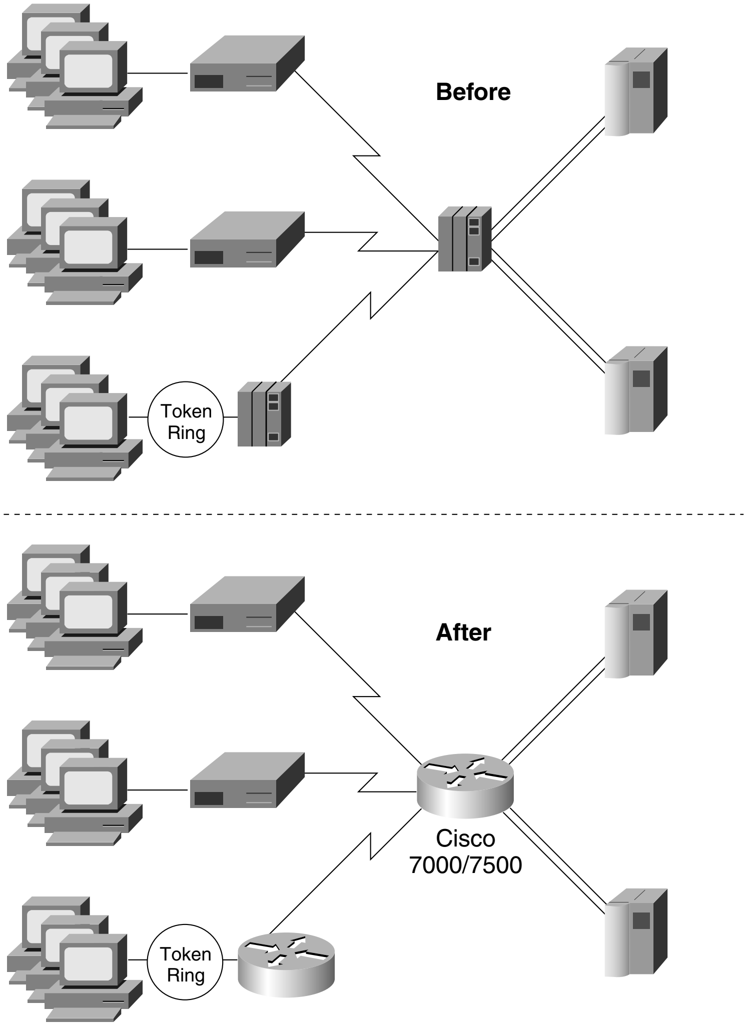

Reasons for Change

The primary reasons for change were to minimize costs and increase throughput and flexibility. The remote IBM 3745 was replaced with a lower-cost Cisco 4500 Series router to eliminate recurring NCP and maintenance charges, consolidate multiprotocol and SNA WAN traffic, and simplify network configuration. The central site FEP was replaced with a channel-attached router to increase channel throughput and to enable TCP/IP on the mainframe in the future.

Design Choices

This enterprise chose not to implement APPN despite having multiple mainframes. The reason is that all SNA sessions were in the same domain. The VTAM in the second mainframe was used just for testing and backup. They decided against implementing two channel-attached routers for redundancy, but did use two CMCCs in a single channel-attached router. This created higher availability than they had previously and provided an option to separate CMCC functionality across multiple CMCCs in the future. They plan eventually to add TN3270 Server capability to the CMCC to allow access to VTAM applications from Web-based clients. They also anticipate a need for TCP/IP on the mainframe. Figure 6-9 shows the logical configuration.

Figure 6-9 Dual CIPs in a Single Router

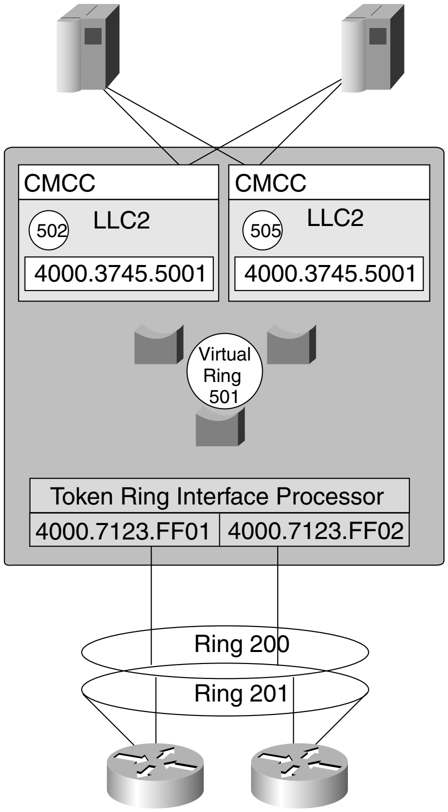

Scenario 4—Combining SNASw with DLSw+

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 Gigabit Ethernet in the data center and the applications are being rewritten to run TCP/IP natively. However, it will be several years before that migration is complete, and in the interim, the customer wants the high availability and design simplicity afforded by having an all-IP data center.

Reasons for Change

This enterprise already uses DLSw+ to transport SNA traffic over an IP backbone. The customer chose not to make an additional investment in SNASw for the branch because the DLSw+ network has been very stable, and if outages occur, they affect only a small portion of the network (by design) and are recovered automatically. However, the customer wants to ensure that a CMCC or channel outage (which today would bring down almost the entire network) can be handled transparently. Hence, the customer is adding SNASw to the SNA routers. The SNA routers use DLSw+ to transport SNA traffic to and from the branch routers and use HPR over IP to transport SNA traffic to and from VTAM. By using HPR over IP directly to VTAM, the customer eliminates any potential looping 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. Finally, this design positions the customer to use a Gigabit Ethernet OSA-Express for SNA traffic.

Design Choices

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+ (to handle a 1000-branch network), including one DLSw+ router used only for backup. Figure 6-10 shows the basic components of this design.

Figure 6-10 Combined SNASw and DLSw+ Design

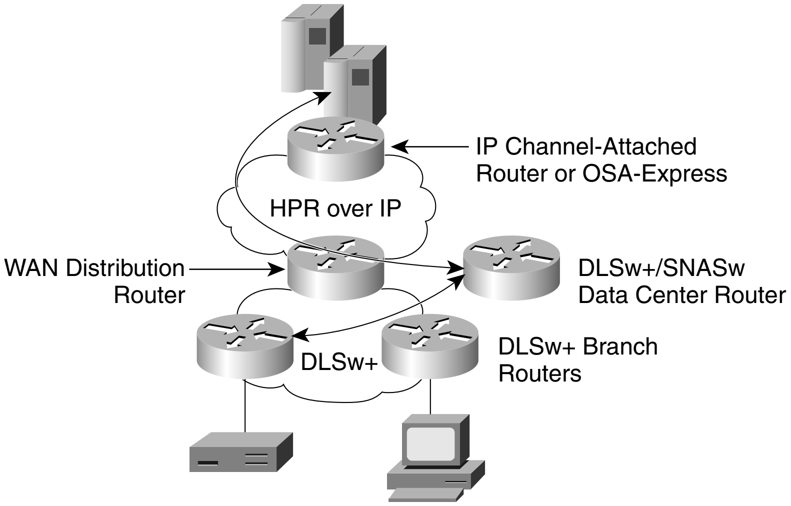

DLSw+/SNASw Data Center Router Configuration

The following configuration is for the SNASw router named Coppito:

```
sh run
Building configuration...
Current configuration:
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname COPPITO
!
ip subnet-zero
ip ftp source-interface TokenRing1/0
ip ftp username cse
ip ftp password csecse
no ip domain-lookup
ip host redclay2 172.18.125.3
!
lane client flush
cns event-service server
!
source-bridge ring-group 400
dlsw local-peer peer-id 10.10.10.99
dlsw remote-peer 0 tcp 10.10.10.1
!
interface FastEthernet0/0
 no ip address
 shutdown
 half-duplex
!
interface TokenRing1/0
 ip address 10.17.1.67 255.255.255.0
 ring-speed 16
!
interface TokenRing1/1
 no ip address
 shutdown
 ring-speed 16
!
interface TokenRing1/2
 no ip address
 shutdown
 ring-speed 16
!
interface TokenRing1/3
 no ip address
 ring-speed 16
 source-bridge 100 1 400
!
interface FastEthernet3/0
 no ip address
 shutdown
 half-duplex
!
interface Ethernet5/0
 ip address 10.10.10.99 255.255.255.0
!
```

```
interface E
thernet5/1
 no ip address
 shutdown
!
interface Ethernet5/2
 no ip address
 shutdown
!
interface Ethernet5/3
 no ip address
 shutdown
!
interface Ethernet5/4
 no ip address
 shutdown
!
interface Ethernet5/5
 no ip address
 shutdown
!
interface Ethernet5/6
 no ip address
 shutdown
!
interface Ethernet5/7
 no ip address
 shutdown
!
snasw pdlog exception file ftp://172.18.125.3/snaswpd1.log
snasw dlctrace file ftp://172.18.125.3/dlctrace1.log
snasw cpname NETA.COPPITO
snasw dlus NETA.MVSD
snasw port HPRIP TokenRing1/3 vnname NETA.EEJEB
snasw port DOWNST vdlc 400 mac 4000.eeee.0000 sap 0x08 conntype nohpr
snasw link TOMVSD port HPRIP ip-dest 172.18.51.1
!
ip classless
ip route 172.18.125.3 255.255.255.255 10.17.1.1
no ip http server
!
line con 0
 exec-timeout 0 0
 transport input none
line aux 0
line vty 0 4
!
end
```

IP Channel-Attached Router Configuration

```
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 ***********************************
```

Scenario 5—Migrating to SNASw only

In this case study, the enterprise demands the highest availability for its SNA applications.

Reasons for Change

The customer has invested a great deal in developing SNA LU 6.2 applications over the years and wants to continue to leverage that investment. The customer has been running separate networks for SNA and IP and has decided to consolidate using SNASw with HPR over IP. The network is already at the latest operating system level and is running APPN in VTAM.

The OSA-Express Gigabit Ethernet card is for TCP/IP environments only. This card supports SNA traffic when SNA is encapsulated in IP using the EE support in OS/390 Version 2, Release 7 or higher.

Design Choices

The customer has 200 regional offices that will run SNASw. From the branch into the S/390, the SNA traffic is transported in IP. Hence, there is no need for SNA routers in the data center. The customer leverages the Cisco IOS QoS features to ensure that the interactive SNA and Telnet traffic take precedence over SNA batch and FTP traffic. Figure 6-11 shows this design.

Figure 6-11 SNASw Design

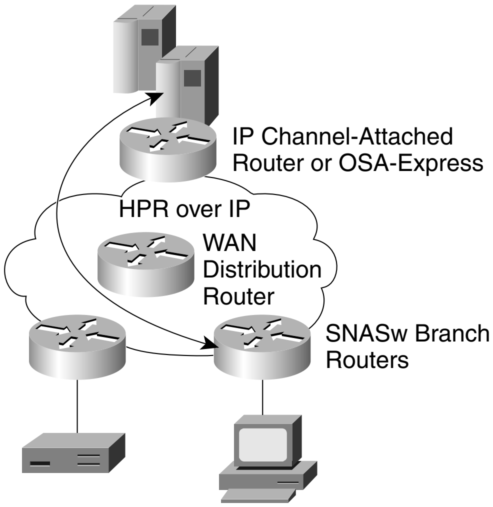

SNASw Branch Router Configuration

```
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#
```

IP Channel-Attached Router Configuration

```
Current configuration:
 !
 version 12.0

 hostname CIPRouter
 !
 enable password lab
 !
 microcode CIP flash slot0:cip216-30
 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 Channel1/0
  no ip address
  no ip directed-broadcast
  no keepalive
 !
 interface Channel1/1
  no ip address
  no ip directed-broadcast
  no keepalive
  cmpc E160 92 EETGJEB READ
  cmpc E160 93 EETGJEB WRITE
 !
 interface Channel1/2
  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=                   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 *******************
```

Scenario 6—Migrating to TCP/IP across CLAW

In this scenario, a customer wants to increase the reliability and robustness of the network by migrating to TCP/IP.

Reasons for Change

Many companies implement a client/server environment by using the database applications available on distributed UNIX or Windows NT servers. Companies also leverage the centralized nature of mainframes to back up this distributed data. IBM's backup product, Tivoli Storage Manager, previously known as ADSTAR Distributed Storage Manager (ADSM), is used to store large amounts of data from distributed platforms to a central mainframe resource (either direct access storage device [DASD] or tape).

Figure 6-12 shows the schematic for this scenario, in which four Sun servers are used for distributed database applications. Each night, the servers use Tivoli Storage Manager to transfer 250 GB data to the centralized mainframe for backup during a three-hour window.

Figure 6-12 Bulk Data Transfer from Distributed UNIX Servers to Central Mainframe

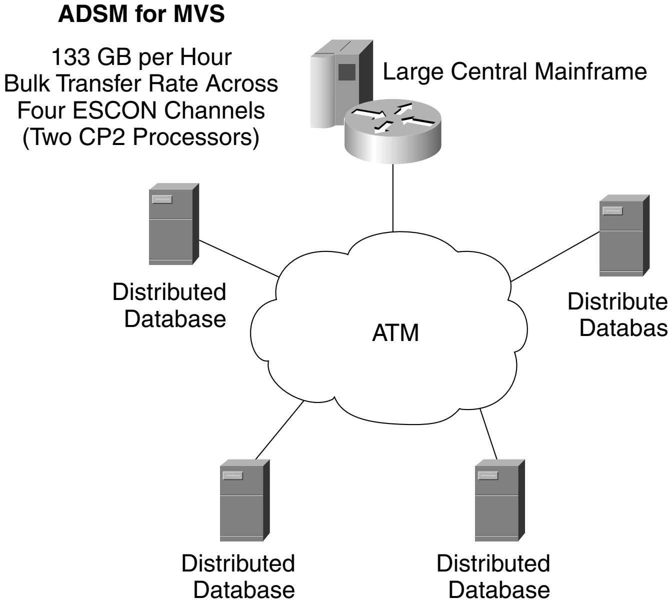

Testing of a CIP in IP Datagram mode determined that a single CIP processor can transfer 18.4 MBps across two ESCON channels. Therefore, two CIP processors can transfer 36.8 MBps. In one hour, the data center router can transfer 133 GB per hour:

36.8 MB per second x 60 seconds per minute x 60 minutes per hour = 133 GB per hour

Therefore, a Cisco 7507 with two CIP cards with dual ESCON interfaces (four ESCON channels) and two Asynchronous Transfer Mode (ATM) interface processors is capable of transferring 133 GB per hour. To determine the amount of time required to transfer the 250 GB of data in the bulk data transfer application example:

250 GB / 133 GB per hour = 1.88 hours, or 112 minutes

As these calculations demonstrate, the Cisco data center router can support the required data transfer rate.

Design Choices

The customer considered several factors before choosing the appropriate components to implement this solution. Although speed and cost were certainly important, the overriding concerns were robustness and reliability. For these reasons, the customer chose CLAW as the channel protocol, because it has been implemented in thousands of data centers and been in widespread use for more than five years.

If your OS/390 host environment supports the use of the Gigabit Ethernet OSA-Express, you should consider the use of OSA-Express with the Tivoli Storage Manager. This solution is optimized to provide very high throughput for bulk data transfer using Large Format Ethernet Frames (also known as Jumbo Frames) and can achieve data transfer rates approaching Gigabit Ethernet speed.

Router Configuration

For configuration examples, see www.cisco.com/warp/public/650/8.html.

Scenario 7—Migrating to TCP/IP across CMPC+

This scenario describes a customer who wants to redesign the OS/390-based data center. This customer wants to migrate from SNA to pure IP using CMPC+.

Reasons for Change

The network architecture group needed to redesign its OS/390-based data center to be the core of a fully enabled e-business and multiservice environment. They had been using CMCC technology for several years to connect both TCP/IP and SNA clients to the S/390s using CLAW channel protocol for the IP traffic and using CSNA for the SNA traffic.

The customer carefully considered and decided that the requirement for successfully moving forward was the ability to control QoS across all applications, from traditional SNA through voice over IP. The customer realized that achieving acceptable QoS would be impossible without removing the protocols that depend on OSI Layer 2 mechanisms for flow control, and that moving to a purely IP transport-based solution would be the best way to optimize positioning for the future.

Design Choices

To reach the goal of building an IP-based backbone network, the customer needed to find a way to transport significant amounts of SNA traffic without depending on the traditional Layer 2-based protocols. EE, which transports SNA data directly over IP, provided the answer. Because this group had extensive experience with CMCC technology, the decision was then largely a matter of deciding which of the IP-capable channel protocols to choose. They decided that CMPC+ provided the best balance of performance for the resulting mix of interactive, batch, and streaming traffic.

Router Configuration

For configuration examples, see www.cisco.com/warp/public/650/8.html.

CMPC+ with TCP/IP Stack Example

This example demonstrates the TCP/IP link for CMPC+ between a host and a Cisco router with a CMCC adapter. The following configuration is for the CIP in the Cisco 7500 Series router:

```
hostname ipclust1
!
microcode CIP flash slot0:cip27-0
microcode reload
!
interface Channel0/1
no ip address
no keepalive
cmpc 0170 00 TG00 READ
cmpc 0170 01 TG00 WRITE
!
interface Channel0/2
ip address 80.12.165.1 255.255.255.0
no ip redirects
no ip directed-broadcast
ip route-cache same-interface
no ip mroute-cache
load-interval 30
no keepalive

tg TG00     ip 80.12.165.2 80.12.165.1
```

In this configuration, the CMPC+ configuration is for the TCP/IP stack on the host. The host IP address of 80.12.165.2 in the transmission group statement corresponds to the IP address for the TCP/IP stack in the TCP/IP profile on the host. The IP address for the CIP is 80.12.165.2.

TCP/IP Profile

The following example shows the TCP/IP profile on the host:

```
ARPAGE 5
telnetparms timemark 600 port 23 dbcstransform endtelnetparms
ASSORTEDPARMS NOFWD ENDASSORTEDPARMS
;
DEVICE mpc4b00  MPCPTP
LINK MPCPLNK2 MPCPTP mpc4b00
;
AUTOLOG
  OEFTPE3
ENDAUTOLOG
INCLUDE TODD.MPCP.TCPIP.PROFILES(PORTS)
HOME
  80.12.165.2   MPCPLNK2
GATEWAY
; NETWORK    FIRST     DRIVER    PACKET   SUBNet mask    subnet value
;             HOP                 SIZE
 80.12.165.1   =    mpcplnk2 4468host
DEFAULTNET 80.12.165.1 mpcplnk244680
BEGINVTAM
    ; Define logon mode tables to be the defaults shipped with the latest
    ; level of VTAM
  3278-3-E NSX32703 ; 32 line screen - default of NSX32702 is 24 line screen
  3279-3-E NSX32703 ; 32 line screen - default of NSX32702 is 24 line screen
  3278-4-E NSX32704 ; 48 line screen - default of NSX32702 is 24 line screen
  3279-4-E NSX32704 ; 48 line screen - default of NSX32702 is 24 line screen
  3278-5-E NSX32705 ; 132 column screen - default of NSX32702 is 80 columns
  3279-5-E NSX32705 ; 132 column screen - default of NSX32702 is 80 columns
    ; Define the LUs to be used for general users
  DEFAULTAPPL ECHOMVSE
; DEFAULTAPPL ECHOMVSE 10.10.1.188
; DEFAULTAPPL NETTMVSE
  DEFAULTLUS
      TCPE0000..TCPE9999
  ENDDEFAULTLUS
  ALLOWAPPL * ; Allow all applications that have not been previously
              ; specified to be accessed
ENDVTAM
DATASETPREFIX TODD.MPCP
start mpc4b00
```

In this TCP/IP profile, the DEVICE specifies the VTAM TRLE mpc4b00 and LINK specifies the link name (MPCPLNK2) associated with the IP address (80.12.165.2) for that link. The host IP address 80.12.165.2 that is specified for the transmission group in the router configuration must be identical to the IP address specified for the transmission group in the router configuration.

TRL Major Node Example

The following configuration shows the TRL major node example:

```
TRL4B00 VBUILD TYPE=TRL
MPC4B00  TRLE  LNCTL=MPC,MAXBFRU=16,X
               READ=(4B00),X
               WRITE=(4B01)
```

In this TRL major node example, the parameter MPC4B00 must be identical to the LINK parameter in the TCP/IP profile.

CHAPTER 7

Using CMCC Network Management Tools

This chapter describes the tools available to manage Cisco channel-attached routers, CMCCs, and Cisco router and switch networks used in your data center solutions. It describes the main network management tools and compares these tools to those commonly used to manage an NCP.

This chapter includes the following information:

- CiscoWorks Blue suite of products
- Comparing a channel-attached router equipped with a CIP to an IBM 3745 with the NCP
- Configuring the router for host management

CiscoWorks Blue Suite of Products

CiscoWorks Blue is a suite of network management products that support management of integrated SNA and router networks. The key products in this suite include:

- Internetwork Status Monitor (ISM)
- CiscoWorks Blue SNA View
- CiscoWorks Blue Maps

These tools speed problem identification, simplify problem isolation, and enable trend analysis. The tools simplify event correlation by consolidating the SNA perspective and the router perspective onto a single console of your choice. In addition to these tools, you can use **show** commands to query CMCC traffic statistics, memory utilization, and cycle utilization. A summary of each product follows.

Internetwork Status Monitor

CiscoWorks Blue ISM for the S/390 provides management and visibility of Cisco devices from the mainframe. Cisco has provided mainframe management of routers since 1996. Powerful features, including a new Web interface, support of Simple Network Management Protocol (SNMP), monitoring of Cisco TN3270 Servers, systems management facility (SMF) logging of events and statistics, and several other usability enhancements provide mainframe operators with the full visibility of Cisco routers, switches, interfaces, and CMCCs from a single mainframe console. By enhancing the NetView management platform, ISM leverages investments in mainframe networks and systems management and provides a reliable, scalable solution for managing Cisco devices.

Note: ISM Version 2.0 works only with IBM's Tivoli NetView for OS/390 (NetView). Earlier releases of ISM also worked with Sterling's SOLVE:Netmaster, which is now Computer Associates' NetworkIT NetMaster. This document discusses the functionality of the latest release, ISM Version 2.0. However, much of the information also applies to the earlier releases.

Protecting Your Investment

ISM gives the data center visibility into the Cisco network while protecting investments in mainframe software and management skills. ISM eliminates the need to use TCP/IP on the host for management, as well as the need to purchase UNIX- or Windows-based software for basic router management. Because ISM uses many of the NetView functions to deliver router management, no retraining or SNMP knowledge is required. Familiar operator consoles with common function keys, help panels, event displays, and network logs help traditional MVS network operators ease their way into distributed network management. For enhanced management, ISM provides SNMP capability in addition to the use of RUNCMDs.

Providing Proactive Network Management

ISM logs a variety of key data for historical and trend analysis, allowing network administrators to manage the network with:

- CMCC performance monitoring data
- Router and interface statistics
- Data from monitoring of Cisco routers, switches, TN3270 Servers, and LocalDirector
- Interface load statistics
- Resource and interface performance monitoring data
- Configuration archiving

By collecting performance data from the network and analyzing historical trends, a network manager can often uncover and avoid problems before they occur. Alerts can be created when CPU or memory thresholds are exceeded on routers and CMCCs.

Providing Increased Productivity and High Availability

Quick, efficient problem resolution translates into higher productivity from network operators and higher availability of your network. ISM provides an array of features that enable operators to rapidly identify, diagnose, and correct problems within the network, including:

- Status-at-a-glance displays for Cisco routers, switches, CMCCs, TN3270 Servers, and interfaces
- Detailed status displays for quick problem diagnosis
- Web interface in addition to 3270 displays
- SMF logging of events and statistics
- Correlation of events
- Integrated command menu for most commonly used commands
- Command-line interface (CLI)
- CMCCs and DSPU management functionality
- Resource grouping to easily facilitate operator's span of control
- Device-specific network management vector transports (NMVTs) for managing Cisco routers
- RIF information displays

Enabling Quick Detection and Correction

Using Web browser or traditional 3270 displays, ISM displays a summary screen of all routers being managed via the service point function or SNMP. Routers are color-coded to indicate their status such as up, down, connect, or performance degraded. Using detailed status displays, operators can quickly diagnose problems by displaying flags next to the troubled router, indicating the nature of the problem. Events forwarded to the mainframe by Cisco resources are correlated with the managed routers, allowing operators to easily select the alerts that apply to a particular router.

Integrating with Mainframe Problem Diagnosis

The ISM management software works in conjunction with the service point feature implemented in the Cisco IOS Software or SNMP agent support implemented in a Cisco device. This combination allows native management of the router from mainframe-based applications and generates NMVTs that are specific to resources and downstream devices. The alerts created in this manner are included in Network Problem Determination Aid (NPDA) displays in NetView. The NMVT alerts use standard code points and require no changes to VTAM or NetView. Automated responses to these alerts can be created in the same manner as any other NMVT generic alert. Figure 7-1 shows the ISM Main Menu Panel in a mainframe environment.

Figure 7-1 ISM Main Menu Panel in a Mainframe Environment

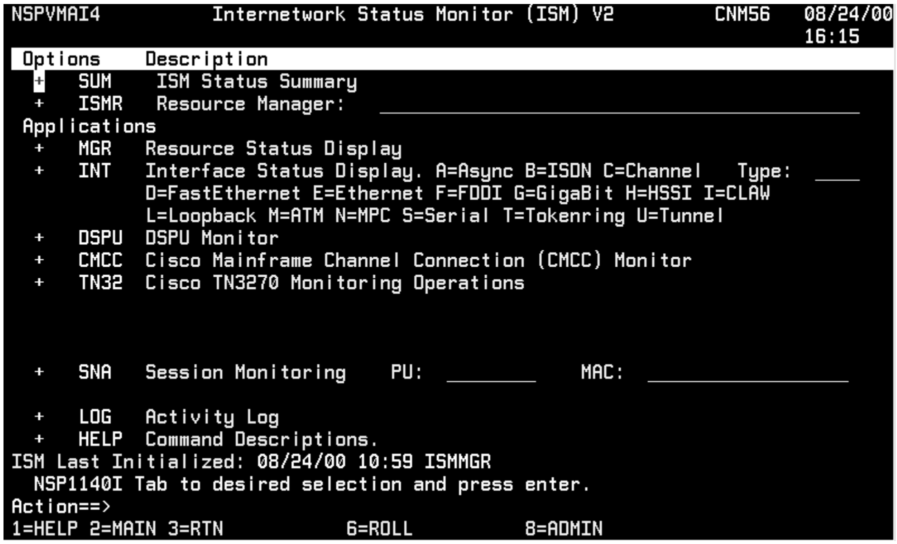

```
NSPVMAI4          Internetwork Status Monitor (ISM) V2        CNM56   08/24/00
                                                                       16:15
 Options    Description
  +   SUM    ISM Status Summary
  +   ISMR   Resource Manager:  ___________________________________________
 Applications
  +   MGR   Resource Status Display
  +   INT   Interface Status Display. A=Async B=ISDN C=Channel   Type: ____
            D=FastEthernet E=Ethernet F=FDDI G=GigaBit H=HSSI I=CLAW
            L=Loopback M=ATM N=MPC S=Serial T=Tokenring U=Tunnel
  +   DSPU  DSPU Monitor
  +   CMCC  Cisco Mainframe Channel Connection (CMCC) Monitor
  +   TN32  Cisco TN3270 Monitoring Operations



  +   SNA   Session Monitoring    PU: ________     MAC: __________________

  +   LOG   Activity Log
  +   HELP  Command Descriptions.
ISM Last Initialized: 08/24/00 10:59 ISMMGR
  NSP1140I Tab to desired selection and press enter.
Action==>
1=HELP 2=MAIN 3=RTN           6=ROLL          8=ADMIN
```

Figure 7-2 shows the Web-based ISM Main Screen, which you can access from your Web browser.

Figure 7-2 Web-Based ISM Main Menu Screen

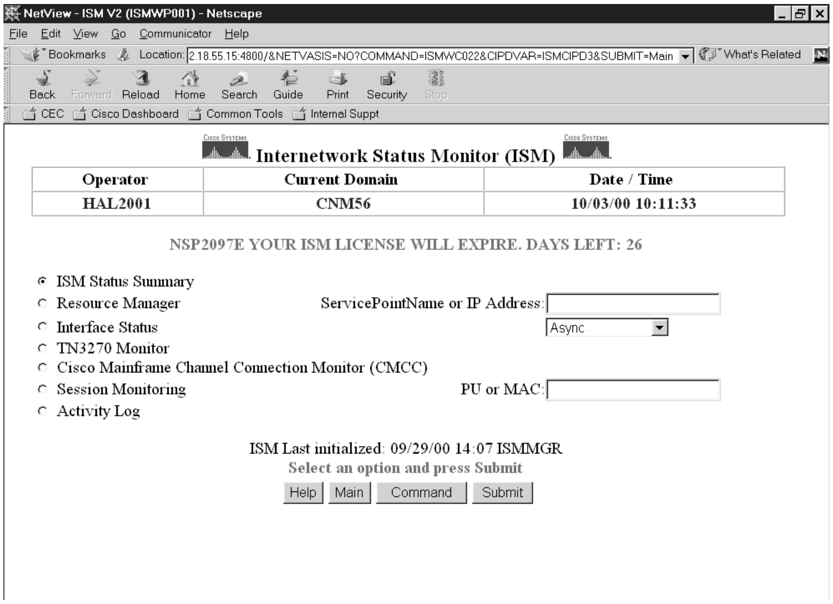
NetView - ISM V2 (ISMWP001) - Netscape
File Edit View Go Communicator Help
Bookmarks
Location: 2.18.55.15:4800/&NETVASIS=NO?COMMAND=ISMWC022&CIPDVAR=ISMCIPD3&SUBMIT=Main
What's Related
Back Forward Reload Home Search Guide Print Security Stop
CEC Cisco Dashboard Common Tools Internal Suppt
Cisco Systems

Internetwork Status Monitor (ISM)

Operator	Current Domain	Date / Time
HAL2001	CNM56	10/03/00 10:11:33

NSP2097E YOUR ISM LICENSE WILL EXPIRE. DAYS LEFT: 26

● ISM Status Summary
○ Resource Manager ServicePointName or IP Address:
○ Interface Status Async
○ TN3270 Monitor
○ Cisco Mainframe Channel Connection Monitor (CMCC)
○ Session Monitoring PU or MAC:
○ Activity Log

ISM Last initialized: 09/29/00 14:07 ISMMGR
Select an option and press Submit

CiscoWorks Blue SNA View

Many organizations are in transition from a native SNA environment (SDLC protocols over slow-speed links or Token Ring) to a mixed SNA/IP environment. Problem determination in this mixed, multiproduct environment can be very difficult with inadequate tools for diagnosis. When an end user calls in with a network problem, it can take a long time to identify the root cause of the problem, which delays resolution.

CiscoWorks Blue SNA View provides an easy-to-use, Web-based interface that takes whatever information an end user can provide and quickly highlights the likely cause of the problem. It integrates with other problem-solving tools such as Tivoli NetView for OS/390, CiscoWorks2000, and Cisco TN3270 Monitor. This integration allows the help desk operator to perform first-level diagnosis on network problems and results in faster problem resolution for your end users.

SNA View extends the capabilities of the CiscoWorks Blue family of network management applications to include correlation and control for the integrated SNA and TCP/IP network. By interacting with the mainframe, SNA View collects PU and LU information and correlates it with information gathered from Cisco devices and Cisco TN3270 Servers. The SNA View operator accesses this information by providing filtering criteria. The filtering criteria can include items such as:

- PU or LU name (or wild-card name)
- IP address (or wild-card address)
- MAC address (or wild-card address)
- Status filter
- Protocol filter (APPN, DLSw, TN3270, and RSRB)
- PU 4 name

SNA View searches its database of SNA sessions and provides a list of sessions that match the filter criteria. The operator can select any of these sessions to receive a graphical end-to-end Session Connectivity Display.

The Session Connectivity Display shows each of the devices that participates in the end user's session and provides status information for each of the devices. The operator can select any of the devices in the session and can optionally "hotlink" to other network management tools depending on the type of device. For SNA resources, such as FEPs or links, the operator can hotlink to the HTML version of Tivoli NetView for OS/390. For Cisco devices, the operator can hotlink to CiscoWorks2000 or CiscoView. If the session is using DLSw, APPN, or RSRB protocols, the UNIX operator can hotlink into the Web-based versions of CiscoWorks Blue Maps. If the session is using TN3270, the UNIX operator can access the Cisco TN3270 Monitor application. Figure 7-3 shows an example of a Session Connectivity Display.

Figure 7-3 Session Connectivity Display

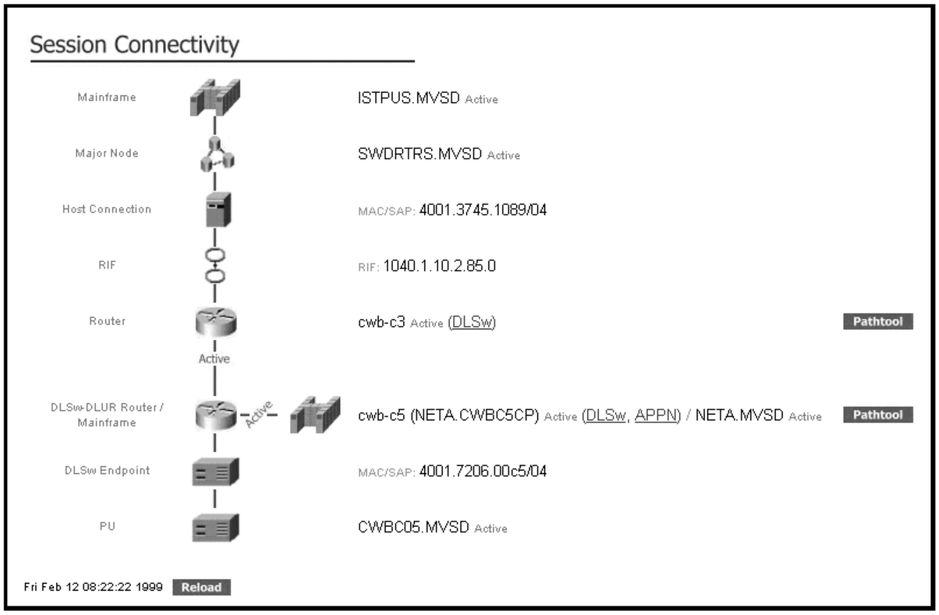

Session Connectivity

Mainframe	ISTPUS.MVSD Active	
Major Node	SWDRTRS.MVSD Active	
Host Connection	MAC/SAP: 4001.3745.1089/04	
RIF	RIF: 1040.1.10.2.85.0	
Router	cwb-c3 Active (DLSw)	Pathtool
DLSw-DLUR Router / Mainframe	cwb-c5 (NETA.CWBC5CP) Active (DLSw, APPN) / NETA.MVSD Active	Pathtool
DLSw Endpoint	MAC/SAP: 4001.7206.00c5/04	
PU	CWBC05.MVSD Active	

Fri Feb 12 08:22:22 1999 Reload

CiscoWorks Blue Maps

CiscoWorks Blue Maps monitors the physical and logical relationships between Cisco routers that support SNA protocols. Maps shows you the status of your combined SNA and IP network through the use of UNIX-based topographical maps and Web-based displays. These displays allow your help desk operators and network administrators to immediately identify problems and to begin problem resolution before end users realize that there is a problem. In addition, Maps correlates information gathered from Cisco routers with information gathered from one or more VTAM domains, enabling the management of PU and LU sessions from within the graphical displays.

Network managers need the right tools for the right job, and troubleshooting a problem within DLSw, APPN, or RSRB networks can be a complex task without those tools. First, you would need to identify the connectivity of a PU or LU to the router supporting that device. Then, you would telnet into that router and issue a series of CLI commands in order to display the circuits and sessions that are passing through the router. After that, you would write the results down, or try to remember them as you telnet from the original router to the router at the other end of the connection. After telneting into the peering router, you would need to issue another set of CLI commands (remembering the correct parameters) and process the results of those commands in order to draw a mental map of the connectivity. Finally, when you had determined the correct path for the connectivity of the session, you would gather data for problem resolution, if you could remember the format of the commands and could interpret the results of those commands. This is quite a laborious process.

Cisco has made all of this easier for you with Maps. It simplifies the process of managing DLSw, APPN, and RSRB networks by monitoring the physical and logical relationships between the Cisco and non-Cisco routers that are supporting these SNA protocols. Maps is an easy-to-use, cost-effective solution that automatically discovers the DLSw, APPN, or RSRB routers and uses information in Management Information Bases (MIBs) to draw a topology map of the protocols. This map shows the physical and logical rings that connect the devices, and the routers are color-coded according to the health of the SNA protocol, providing status information at a glance. Additional information is just a mouse-click away, providing quick access to detailed status of peer connections, traffic statistics, and error statistics. No longer do you need to remember protocol-specific CLI commands or learn MIB values.

Cisco has provided seamless integration of Maps with CiscoView and the Path Tool application, providing the ability to gather detailed information about the health of the device, such as the status of interfaces, ports, and the IP connectivity between resources. Maps goes that extra step by correlating PU and LU information from VTAM to provide end-to-end dependency views, showing how your end user's sessions traverse your IP environment—an essential picture for diagnosing problems.

As network managers are moving rapidly to support Web technologies and Internet standards as the basis for the next generation of enterprise management solutions, Maps has embraced the "management intranet" by providing a Web-based interface. Maps can be integrated into the CiscoWorks2000 framework, providing a "portal" of management applications utilizing the power of Web technologies in solving management problems.

Comparing a Channel-Attached Router Equipped with a CIP to an IBM 3745 with the NCP

This section compares the network management tools that IBM provides for the NCP with those available for a similar CMCC/router configuration. The following functions are examined:

- Alerts
- Statistics
- Console support
- Trace/debug
- Connectivity test
- Memory display/dump
- Recovery
- Performance monitoring
- Configuration management
- Router configuration for host management

Generating Alerts

Both the NCP and the CMCC (via the service point) provide alerts when a resource fails.

Alerts are either generated by or passed through the NCP to VTAM. In VTAM, the alerts are forwarded to NetView. The NCP generates alerts for all the resources it manages and forwards the alerts it receives from an external resource. The NCP does not create traps or support SNMP.

The CMCC/router provides alerts for SNA resources and convert some traps to alerts. It also creates traps for all the alerts. Table 7-1 compares the NCP and CMCC alerts. In the case of switched SNA resources operating on a router SDLC interface, alerts are generated for all SDLC-related errors.

Table 7-1 Comparison of Alert Support for the NCP and the CMCC

Alert Support Comparison	NCP	CMCC
Local Interface Failure	Yes	Yes
Downstream Alerts	Yes	Limited
Resolved Alerts	No	Possible (ISM with Syslog)
Threshold	Slowdown	Possible (ISM)

When VTAM initiates intensive-mode recording (IMR), the NCP can generate alerts for soft errors. Figure 7-4 shows an example of an alert received from a router.

Figure 7-4 NetView Alert Screen

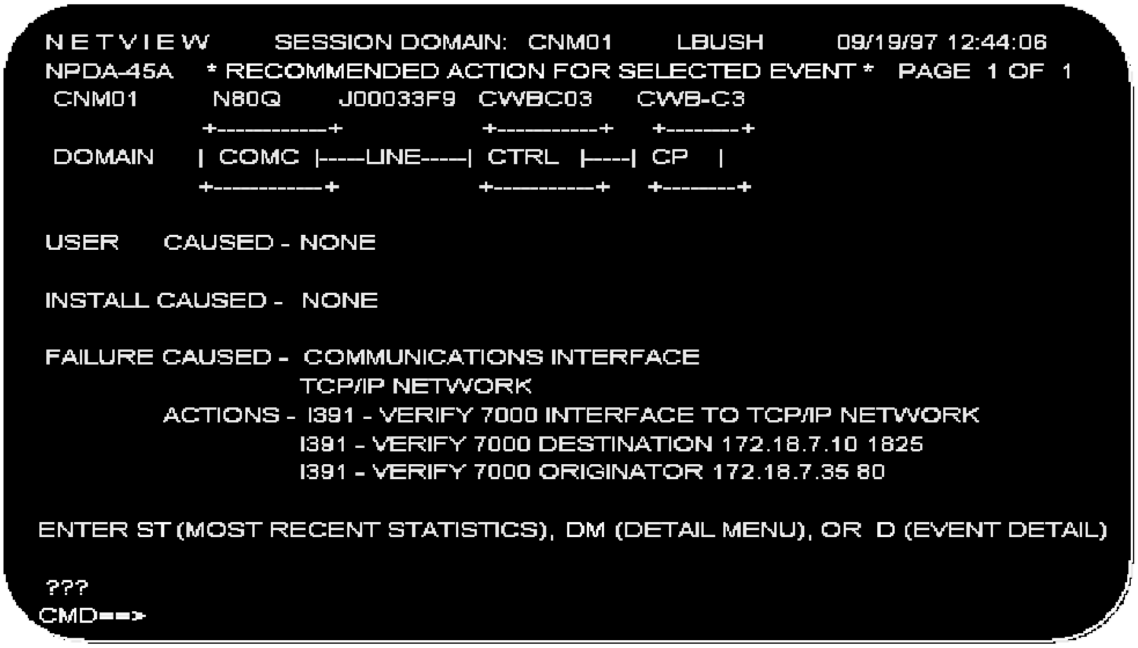

```
NETVIEW       SESSION DOMAIN:  CNM01       LBUSH        09/19/97 12:44:06
NPDA-45A    * RECOMMENDED ACTION FOR SELECTED EVENT *   PAGE  1 OF  1
 CNM01        N80Q       J00033F9  CWBC03     CWB-C3
             +----------+           +----------+    +-------+
 DOMAIN      | COMC  |----LINE----| CTRL  |----| CP    |
             +----------+           +----------+    +-------+

USER      CAUSED - NONE

INSTALL CAUSED -  NONE

FAILURE CAUSED -  COMMUNICATIONS INTERFACE
                  TCP/IP NETWORK
         ACTIONS - I391 - VERIFY 7000 INTERFACE TO TCP/IP NETWORK
                  I391 - VERIFY 7000 DESTINATION 172.18.7.10 1825
                  I391 - VERIFY 7000 ORIGINATOR 172.18.7.35 80

ENTER ST (MOST RECENT STATISTICS), DM (DETAIL MENU), OR  D (EVENT DETAIL)

 ???
CMD==>
```

Collecting Statistics

NCP provides statistics for all links, lines, and PUs. NetView has the ability to log the statistics to SMF. NCP provides statistics for the following conditions:

- When the NCP is shut down normally
- When a resource fails or is made inactive
- When a counter that relates to the interface is filled, such as traffic or soft errors

Table 7-2 summarizes the statistical features of the NCP and the CMCC.

Table 7-2 Statistics Summary

Statistics Summary	NCP	CMCC
End of Day (EOD)	Yes	No
Threshold	Yes	No
Solicited	No	Yes (ISM)
Archived	Yes (NPDA)	Yes (ISM)

Figure 7-5 shows an example of a statistical record for an NCP managed resource.

Figure 7-5 Statistical Record for an NCP-Managed Resource

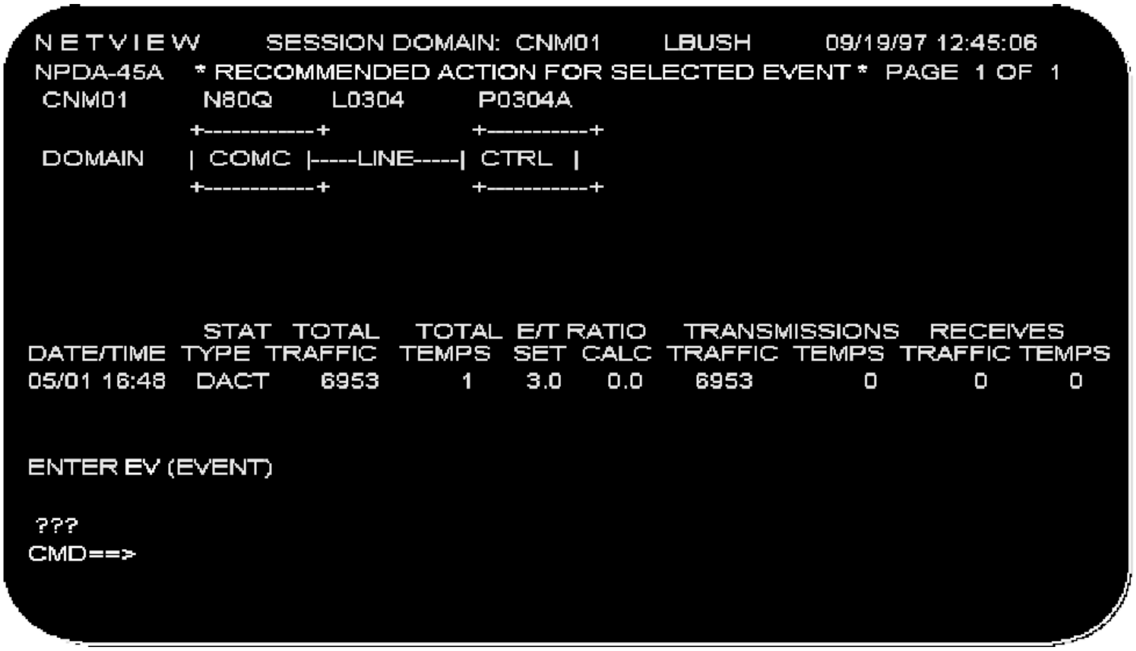

```
NETVIEW       SESSION DOMAIN: CNM01      LBUSH        09/19/97 12:45:06
NPDA-45A    * RECOMMENDED ACTION FOR SELECTED EVENT *  PAGE  1 OF  1
 CNM01        N80Q       L0304        P0304A
             +-----------+            +----------+
 DOMAIN      | COMC |-----LINE-----| CTRL  |
             +-----------+            +----------+




              STAT  TOTAL    TOTAL  E/T RATIO    TRANSMISSIONS    RECEIVES
DATE/TIME  TYPE  TRAFFIC   TEMPS   SET  CALC  TRAFFIC  TEMPS  TRAFFIC TEMPS
05/01 16:48   DACT     6953          1      3.0     0.0     6953           0          0         0


ENTER EV (EVENT)

 ???
CMD==>
```

ISM can collect statistics based on user-defined intervals. Also, ISM can collect and archive router CPU and memory utilization data. Interface statistics can be collected and archived. ISM monitors the CIP, CPU, and memory, and monitors and archives channel statistics. Figure 7-6 shows an example of records collected for an ISM-managed interface.

Figure 7-6 Web-Based Statistical Record Collected for an ISM-Managed Interface

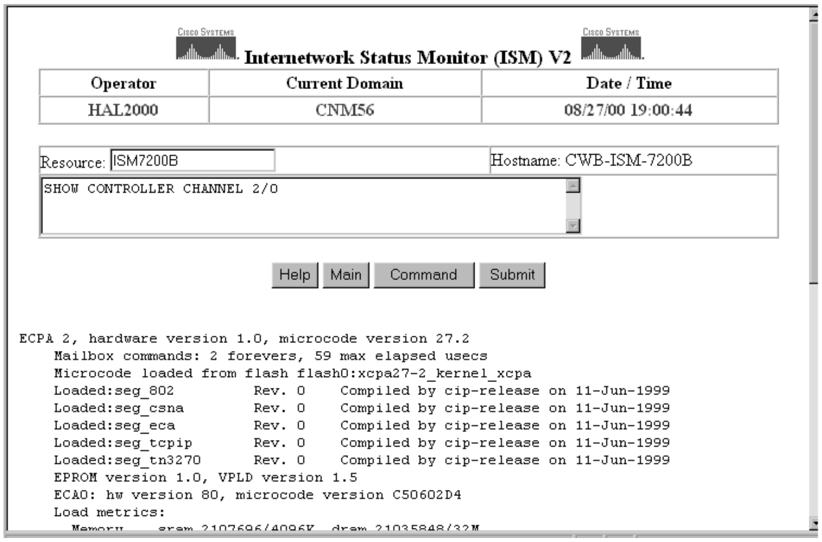

Operator	Current Domain	Date / Time
HAL2000	CNM56	08/27/00 19:00:44

Resource: ISM7200B   Hostname: CWB-ISM-7200B

SHOW CONTROLLER CHANNEL 2/0

```
ECPA 2, hardware version 1.0, microcode version 27.2
    Mailbox commands: 2 forevers, 59 max elapsed usecs
    Microcode loaded from flash flash0:xcpa27-2_kernel_xcpa
    Loaded:seg_802         Rev. 0    Compiled by cip-release on 11-Jun-1999
    Loaded:seg_csna        Rev. 0    Compiled by cip-release on 11-Jun-1999
    Loaded:seg_eca         Rev. 0    Compiled by cip-release on 11-Jun-1999
    Loaded:seg_tcpip       Rev. 0    Compiled by cip-release on 11-Jun-1999
    Loaded:seg_tn3270      Rev. 0    Compiled by cip-release on 11-Jun-1999
    EPROM version 1.0, VPLD version 1.5
    ECA0: hw version 80, microcode version C50602D4
    Load metrics:
      Memory    sram 2107696/4096K  dram 21035848/32M
```

Providing Console Support

The FEP allows attachment of a console used for service support of the FEP. The console (MOSS) has limited functionality with the FEP and is not supported by NCP. Console users must have technical expertise because the console is primarily used to display and alter storage on the FEP.

The CMCC/router offers various options:

- A terminal can be attached to the router console interface
- A terminal can be attached via a Telnet session
- NetView can be attached as a console when the service point function is implemented in the router

The console support for the router allows operators to perform router functions, as listed in Table 7-3. The MOSS console is usually located next to the FEP, and access to the console is usually restricted. The console support for the router is local or remote and has security features to control access and command level.

Table 7-3 Console Support

Console Support	NCP	CIP/Router
Console	Yes (MOSS)	Yes
Telnet Access	No	Yes
NetView Access	Yes	Yes (RUNCMD)

Providing Trace and Debug Facilities

VTAM provides trace facilities that you can use to trace either the NCP or the CMCC. NCP has a line trace that is initiated from VTAM. The output is sent back to VTAM and recorded using the Generalized Trace Facility (GTF). A VTAM buffer trace is available for any resource that is known to VTAM. Table 7-4 contrasts the trace/debug support of the NCP and CMCC/router.

Table 7-4 Trace/Debug Facilities for the NCP and CMCC/Router

Trace/Debug	NCP	CMCC/Router
Data Trace	Yes (VTAM)	Yes (VTAM)
Line Trace	Yes (VTAM)	No (See debug)
Packet Trace	No	Yes (DEBUG)

The router provides debug facilities that allow detailed problem determination. In most cases, the debug facility requires an authorized operator. Cisco recommends that you perform the trace type of debug operations through a Telnet session rather than across the service point interface.

The Cisco IOS Software provides extensive debug facilities. For more detail, see the documentation under the appropriate Cisco IOS release at www.cisco.com/univercd/cc/td/doc/product/software/index.htm.

Performing Connectivity Tests

The connectivity test allows you to verify a remote connection. In SNA networks, NPDA provides functions that allow you to verify a remote connection. The LPDA function was added to support modems that have the LPDA feature.

The router allows you to ping a remote resource, if it has an IP address. IPM can perform connectivity tests and report exceptions. You can also create connectivity features in ISM that direct specific routers to perform a connectivity test.

Displaying and Dumping Memory

VTAM provides two types of facilities to obtain memory information from the NCP, which are discussed in the following sections.

Displaying Storage Information

Use the following command to obtain storage information from the NCP:

```
DISPLAY NET,NCPSTOR,ID=&NCP,ADDR=&ADR,LENGTH=&LEN
```

NetView provides a CLIST (NCPSTOR) to simplify the use of this command:

```
NCPSTOR NCP572P,260
```

The NCP displays the following responses:

```
IST097I NCPSTOR ACCEPTED
 IST244I NCP STORAGE FOR ID = NCP572P
 IST245I 000260 81C2282F 104828AE 415CF00F 991528B3
 IST245I 000270 0108E1F0 80804154 A821D410 25B9F2A0
```

Dumping an Active NCP

Use the following command to dump router memory from the NCP:

```
MODIFY NET,DUMP,&OPTIONS
```

NetView provides a CLIST (NCPDUMP) to simplify the use of the command:

```
NCPDUMP NCP1,DYNA,PDS=NCPDUMP
```

Cisco routers can display router memory statistics. For more detail, see the documentation under the appropriate Cisco IOS release at www.cisco.com/univercd/cc/td/doc/product/software/index.htm.

Recovering the Interface

When an NCP fails or an interface on an NCP fails, you must add automation routines to perform the recovery. Without automation routines, the interface remains inactive and the NCP will not try to recover the interface.

On the Cisco router, interfaces attempt to recover unless they are administratively down. However, you must add automation routines in NetView to recover the channel when the CMCC is reloaded.

Monitoring Performance

Performance monitoring determines whether the performance of the FEP and interface is acceptable. Also, monitoring is performed for planning purposes. In most environments, products such as NETSPY or NPM are used to monitor performance.

With the router, you can use the **show** commands to monitor buffer utilization, memory utilization, or CPU. Figure 7-7 shows an example of router performance data archived by ISM.

Figure 7-7 SM Router CPU/Memory Utilization Data

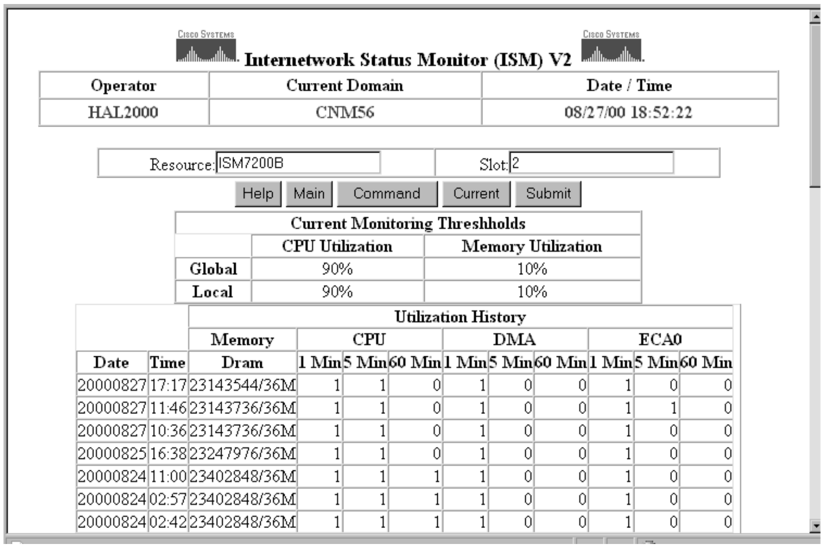

Internetwork Status Monitor (ISM) V2

Operator	Current Domain	Date / Time
HAL2000	CNM56	08/27/00 18:52:22

Resource: ISM7200B

Slot: 2

Current Monitoring Threshholds		
	CPU Utilization	Memory Utilization
Global	90%	10%
Local	90%	10%

		Utilization History									
		Memory	CPU			DMA			ECA0		
Date	Time	Dram	1 Min	5 Min	60 Min	1 Min	5 Min	60 Min	1 Min	5 Min	60 Min
20000827	17:17	23143544/36M	1	1	0	1	0	0	1	0	0
20000827	11:46	23143736/36M	1	1	0	1	0	0	1	1	0
20000827	10:36	23143736/36M	1	1	0	1	0	0	1	0	0
20000825	16:38	23247976/36M	1	1	0	1	0	0	1	0	0
20000824	11:00	23402848/36M	1	1	1	1	0	0	1	0	0
20000824	02:57	23402848/36M	1	1	1	1	0	0	1	0	0
20000824	02:42	23402848/36M	1	1	1	1	0	0	1	0	0

Configuration Management

Except for dynamically adding lines, and PUs and LUs, NCP configurations must be assembled and loaded to add new features. This compilation is performed at the mainframe. A copy of the source is passed to VTAM so that it knows what has been generated in the NCP.

For the CMCC/router environment, VTAM does not need to know what is configured in the router. However, you must define a TYPE=XCA major node that identifies the channel the CMCC is using. All resources connecting to VTAM via the CMCC must be defined in a TYPE=SWNET. NCPs are loaded via the host channel. Routers are loaded from an FTP server.

ISM allows you to archive the router configuration in the mainframe. ISM also has the capability to discover and monitor all interfaces configured in the router.

Configuring the Router for Host Management

XCA Major Node

```
*DDDLU LUGROUP FOR TN3270
*
*DATE CHANGED   WHO       WHAT
*------------   ----     ------------------------------
****************************************************
XCAPUGEN        VBUILD  TYPE=XCA
X31PR04         PORT    MEDIUM=RING,ADAPNO=4,SAPADDR=4,CUADDR=8C0,TIMER=90
X31PR04         PORT    MEDIUM=RING,ADAPNO=4,SAPADDR=4,CUADDR=8C0,TIMER=90, X
                        TGP=TRING16M,VNNAME=NETA.CNNNET1,VNGROUP=CNNGRP1
CNNGRP1         GROUP   DIAL=YES,ISTATUS=ACTIVE,ANSWER=ON,CALL=INOUT,X
AUTOGEN=(100,L,P)
GRP390T5        GROUP   DIAL=NO
LN390T5         LINE    USER=SNA,ISTATUS=ACTIVE
P390T5          PU      MACADDR=400170000390,TGN=1,SAPADDR=04,SUBAREA=39,X
                        PUTYPE=5,ISTATUS=ACTIVE
```

Switched Major Node Definition

```
*SWDRTRS VBUILD TYPE=SWNET
***************************************
* SW MAJ NODE FOR LAB AND RUNCMD TESTING OF ROUTERS
*
* LAB TEST ROUTER  CWBC01
*
* CWBC01  PU  ADDR=01,  X
              PUTYPE=2, X
              IDBLK=05D,X
              IDNUM=CC001,X
              DISCNT=(NO), X
              ISTATUS=ACTIVE,X
              MAXDATA=521,X
              IRETRY=YES, X
              MAXOUT=7, X
              PASSLIM=5,X
              MAXPATH=4Router Configuration (Partial)
Building configuration...
Current configuration:
!
version 11.2
service udp-small-servers
service tcp-small-servers
!
hostname cwb-c1
!
boot system flash slot0:c7000-js-mz
boot system mzallocc/c7000-j-mz 171.69.160.22
 enable password -- suppressed --
!
microcode CIP flash slot0:cip208-0_kernel_hw4
microcode reload
ip subnet-zero
ip domain-name cisco.com
ip name-server 171.69.160.21
ip name-server 171.68.10.70
ip accounting-list 0.0.0.1 255.255.255.0
source-bridge ring-group 900
source-bridge remote-peer 900 tcp 172.18.9.17
source-bridge remote-peer 900 tcp 172.18.9.145
dlsw local-peer peer-id 172.18.9.161 promiscuous
!
> DSPU is required for focal point connection via the CIP.
dspu rsrb 325 1 900 4000.7000.0001
dspu rsrb enable-host lsap 4
!
dspu host CWBC01 xid-snd 05dcc001 rmac 4000.3333.4444 rsap 4 lsap 4 focalpoint
!
dspu rsrb start CWBC01
!
interface Tunnel0
no ip address
!
interface Ethernet1/0
no ip address
shutdown
no mop enabled
!
```

```
interface Ethernet1/1
description ethernet to hub 2
no ip address
shutdown
no mop ena
bled
!
interface Ethernet1/2
no ip address
shutdown
no mop enabled
!
interface Ethernet1/3
no ip address
ip accounting output-packets
ip accounting access-violations
shutdown
> Listing terminated
```

Comparing a CMCC to the IBM 3745/3746, 3172, and OSA

Currently, the key solutions in use to connect your network to an IBM S/390 mainframe are the IBM 3745/3746 FEP, the IBM 3172 Interconnect Controller, the IBM Open Systems Adapters (OSAs), or a Cisco router with a CMCC. A Cisco router can directly replace an IBM 3172 without any loss of function, and it will improve performance and availability. In many cases, you can use a Cisco router with a CMCC as a higher-performance, lower-cost alternative to a FEP. However, some functions might still require a FEP.

This chapter compares the functions provided by provided by a Cisco channel-attached router to the following IBM platforms:

- IBM 3745/3746 FEP
- IBM 3172 Interconnect Controller
- IBM OSA

The chapter ends with a summary comparison of the CMCC and the IBM channel connectivity solutions.

IBM 3745/3746 FEP

Historically, IBM's primary solution for mainframe access has been the FEP. The FEP offers a great deal of functionality for subarea networks and legacy protocols. However, only the largest networks use most of the functionality provided by the FEP; most small networks use only a subset of this functionality. In addition, networks are changing rapidly and the typical enterprise network now supports a multitude of protocols, LANs, WANs, and device types. High-performance substitutes, such as LANs, high-speed serial lines, and Frame Relay have replaced low-speed serial lines. The FEP has not kept up with the requirements of today's enterprise networks so other networking gear is required to augment or replace the FEPs. If you are considering replacing some or all of your FEPs, first determine which functions your FEP is providing today so that you do not lose any of these functions as you move forward to CMCC.

FEPs have the following key functions in today's networks:

- *SNA session routing*—SNA session routing is required in environments with multiple data centers or Advanced Communications Function (ACF)/VTAM application hosts and a high volume of cross-domain SNA traffic. SNA session routing can be important in environments with distributed AS/400s.
- *SNA COS*—SNA COS allows prioritization of SNA traffic between the FEPs and the mainframes and is important in environments with SNA backbones. SNA COS is less important in environments that have consolidated the FEPs in the data center. In this case, either there is no FEP-to-FEP traffic, or the FEPs are connected at the data center over high-speed LANs that do not have bandwidth contention problems. However, some networks take advantage of Link Services Prioritization (LSPRI), which provides transmission priority based on COS for outbound traffic (for example, FEP to cluster controller).

- *Serial line concentration*—FEPs can concentrate large numbers of low-speed (9.6-kbps) serial lines. However, as networks migrate to high-speed WAN backbones, the need for high-density, low-speed serial connectivity decreases.
- *Switched SDLC*—Some enterprises rely on switched SDLC to support transient SNA connections to small branch offices or to provide switched network backup. As SDLC is being replaced by multiprotocol data links, switched SDLC requirements are diminishing. In place of SDLC, protocols such as Integrated Services Digital Network (ISDN), Point-to-Point Protocol (PPP), and Serial Line Interface Protocol (SLIP) are being used to provide multiprotocol or IP-switched line support.
- *SNA boundary network node (BNN) function*—FEPs provide an SNA BNN function, which includes polling, converting from local addresses to SNA addresses, and converting exchange identification (XID). In the absence of remote FEPs, local FEPs can perform these functions. In the absence of any FEPs, ACF/VTAM can perform most of these functions.
- *SNA Network Interconnection (SNI)*—Many enterprises use FEPs for SNI to allow independent SNA networks to communicate. There are other alternatives, such as the SNASw border node function and electronic data exchange over the Internet; however, any change on one side requires a change on the other side, so this migration will be a slow one.
- *SSCP takeover*—With this facility, if an owning VTAM goes down, another VTAM can assume ownership of those resources without disrupting any existing application sessions. The NCP plays a role in allowing this takeover.
- *Extended recovery facility (XRF)*—The XRF is a program that allows one VTAM application to take over for another. The XRF code in the NCP plays a key role in supporting this capability.
- *X.25 support*—X.25 Interconnection allows the NCP to act as an X.25 packet switch. NCP Packet Switching Interface (NPSI) allows the NCP to connect to other resources over X.25 networks. X.25 Interconnection supports both SNA and non-SNA devices. For non-SNA (Asynchronous and Binary Synchronous Communications Protocol) devices, it supports conversion to SNA.
- *Specialized program products that support custom or older applications*—Network Routing Facility (NRF) provides routing inside the NCP without VTAM participation. An emulation program allows the IBM 3745 to connect to Basic Telecommunications Access Method (BTAM) in an IBM mainframe.
- *Legacy protocols*—The FEP supports program products, such as Non-SNA Interconnection (NSI) for Bisynch conversion, Airline Line Control Interconnection (ALCI) for airline line control protocol transport, and Network Terminal Option (NTO) for synchronous conversion. You can install these products in the FEP to handle non-SNA protocols. Legacy protocols are older protocols that are declining in usage.

FEPs such as IBM 3745/3746 hardware support IBM NCP software to provide network control and routing for SNA subarea networks. The IBM 3745/3746 supports high-speed attachments for IP/HPR flows, such as 155–Mbps ATM, Primary Rate Interface (PRI) ISDN, and ESCON MPC+.

Using the Cisco Channel-Attached Router as a FEP Alternative

The Cisco channel-attached router can be used as an alternative to the IBM 3745/3746 FEP. The Cisco router and CMCC combination focuses on the key features that most IBM customers use, such as mainframe channel attachment for both SNA and TCP, SNA routing, SNA COS, and access to SDLC- and LAN-attached resources.

Looking at the key FEP functions identified in the previous section, the Cisco IOS Software and the CMCC offer a way to address most of the key FEP functions while providing a higher-performing, multipurpose channel gateway. For functions not addressed by the CMCC, one or more FEPs still may be required.

SNA Session Routing

The Cisco IOS Software supports native APPN routing through the SNASw feature, and with DLUR it can support native SNA routing for legacy 3270 traffic in an SNASw network. In many environments, SNASw is the logical progression from subarea SNA. SNASw is more dynamic and less labor-intensive to maintain than a subarea network, and it extends the number of SNA network addressable units (NAU) beyond the 64,000 limit per subarea SNA domain.

SNA COS

The SNASw feature also preserves SNA COS for both APPC and legacy 3270 traffic. If the Cisco SNASw feature is installed only in central site DLSw+ routers, it provides outbound prioritization based on COS, similar to LSPRI in the FEP. In a multiprotocol environment, Cisco queuing algorithms such as Custom Queuing can be used to reserve bandwidth for SNA traffic over DLSw+ (DLSw+ supports LU and SAP prioritization). SNASw EE supports SNA COS to IP type of service (ToS) mapping (SNA transmission priority to IP precedence mapping) for both inbound and outbound traffic in a bidirectional fashion between an EE-enabled S/390 host and an SNASw EE router, allowing the service policy agent within CS/390 to enforce QoS policies.

Serial Line Concentration

The Cisco 3600 Series router supports up to 54 serial lines, or 36 serial lines plus one LAN. The Cisco 7200 Series router supports 48 serial lines and one LAN, which can be Fast Ethernet. Either the Cisco 3600 or the Cisco 7200 Series router is a good solution for low-speed SDLC serial line concentration. The MultiChannel Interface Processor (MIP) card is the best solution for high-speed serial concentration when the remote branches have routers and connect over 56- or 64-kbps lines or ISDN Basic Rate Interface (BRI). The Cisco 7500 Series router MIP card supports two channelized T1/E1s or two PRI ISDN lines, supporting up to 48 56-kbps or 64-kbps remote sites per card. You can install multiple MIP cards in a Cisco 7500 Series router. If your current network is pure SNA and your FEPs connect 200 or more low-speed (19.2 kbps or below) serial lines, the branches are too small to justify a router. Although a packet-switched service such as X.25 or Frame Relay is not an option, the FEP still may be the most cost-effective solution.

Switched SDLC

The Cisco IOS Software transports multiprotocol traffic, including SNA, over switched services. However, it does not support dial-out to switched SDLC devices. (Dial-in requires that you code **sdlc role prim-xid-poll** on the appropriate serial interface.)

SNA BNN Functions

The Cisco IOS Software can reduce mainframe cycles by providing several boundary functions such as remote polling, group poll support, and DSPU concentration. Using SNASw and DLUR, Cisco routers can provide many functions provided by a FEP.

Autonomous Network Connection

SNI connections require a FEP in at least one of the connecting networks. The Cisco IOS Software allows connection to an SNI gateway, but it does not provide SNI gateway functionality, as shown in the following examples:

- If the inter-enterprise connection uses back-to-back SNI gateways, at least one FEP is required in each independent SNA network.
- If the inter-enterprise connection can be an adjacent SNI configuration, one network can keep a FEP to provide the SNI gateway function, and the attaching network can replace FEPs with CIPs. The downside to this alternative is that certain topology changes in one network (for example, adding a new subarea node) might require changes in the other network.

- If the inter-enterprise connected hosts are APPN-enabled, they can eliminate SNI connections using APPN border node support (either extended or peripheral). APPN border node allows networks with different NETIDs to establish CP-to-CP sessions with each other (SNASw does not play any role in host-to-host border node connections). Cisco routers and multilayer switches can also provide IP transport between hosts that implement 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 hosts.

Casual connection can be used, eliminating the FEP requirements for one network. This connection supports primary LU (application) initiated sessions only.

SSCP Takeover

The SNASw DLUR feature of the Cisco IOS Software fully supports the SSCP takeover facility.

XRF

This product requires an NCP. There is no channel-attached router equivalent.

X.25 Support

The Cisco IOS Software can be configured as an X.25 packet switch and supports transport of SNA over an X.25 backbone. However, there is no comparable function to provide asynchronous or bisynchronous conversion to SNA. (The CMCC does support the TN3270 Server, which provides conversion from TN3270 to SNA.)

Specialized Program Products that Support Custom or Older Applications

There is no function comparable to NRF in the Cisco IOS Software. There is no feature comparable to the FEP in the Cisco IOS Software.

Legacy Protocols

Although the Cisco IOS Software can duplicate some special protocols supported by the FEP, such as asynchronous and bisynchronous tunneling, comparable protocol support (that is, conversion to SNA) is not provided. (The CMCC does support the TN3270 Server, which provides conversion from TN3270 to SNA.)

Benefits of Using a Cisco 7000 or 7500 Series Router with a CMCC

A Cisco channel-attached router with a CMCC offers additional features that are not available in an IBM 3745. These features include:

- *Multipurpose*—The CMCC provides a state-of-the art, high-performance solution for mainframe connectivity for access not only to SNA applications but to TCP/IP applications as well. As networks begin to offer intranet and Internet services, tying the mainframe into TCP/IP networks with features such as TN3270 Server enables you to leverage your mainframe investment. The NCP's support of TCP/IP is limited.
- *Higher speed*—The CMCC offers a tenfold improvement in performance over an IBM 3745 model 200 for SNA, and even larger for TCP/IP. Many networks have reached capacity for their existing IBM 3745s. Instead of investing more money in older technology, organizations are migrating to multifunction channel solutions.
- *Connectivity*—The FEP has limited connectivity. It does not support Fiber Distributed Data Interface (FDDI), ATM, LANE, Fast Ethernet, Switched Multimegabit Data Service (SMDS), T3, or even Ethernet (for SNA). (The IBM 3746 900 expansion frame supports Ethernet with an imbedded 8229 translational bridge.)
- *Lower cost*—The CMCC in a Cisco 7500 Series router can save your organization money because there is no recurring licensing fee, and the resale value of the IBM 3745 often pays for the CMCC. In leasing environments, the payback period is 18 to 24 months.

In summary, the Cisco 7500 Series router in conjunction with a CMCC offers many benefits such as speed, connectivity, and flexibility to an IBM enterprise network. By minimizing the number of FEPs required in a network, the CMCC offers a means to reduce network costs while improving performance. However, some FEPs still may be required for the following:

- SNI connections to other enterprises or divisions
- Bisynchronous, asynchronous, or ALC conversion
- Specialized functions, such as an emulation program, XRF, or NRF

IBM 3172 Interconnect Controller

The IBM 3172 Interconnect Controller, introduced in the late 1980s, was IBM's premier solution for LAN-to-mainframe connectivity. The IBM 3172 was essentially a rugged PS/2 server that was based on a Micro Channel Architecture (MCA) bus. The IBM 3172 supported both TCP/IP and SNA traffic to the host. It supported a limited number of LAN and WAN connections.

The IBM 3172 is no longer sold by IBM, but there is a large installed base of the devices worldwide. Therefore, this discussion will focus on the reasons to replace existing IBM 3172s with a CMCC solution.

Through software options, the IBM 3172 offers a variety of different capabilities. Its native support is a basic kernel operating environment called Interconnect Controller Program (ICP), which provides functionality similar to that provided by a CMCC operating in IP Datagram mode (albeit with a different channel protocol used). The IBM 3172 offers three other options that require the OS/2 operating system: TCP/IP Offload, SNA Communications Program, and Multiprotocol Extensions. The TCP/IP Offload option is similar to that offered on the CMCC. However, because of the limited processing capability of the platform and the fact that the Offload option is provided on top of a generic OS/2 operating system, the total throughput of this option was always very low. The SNA Communications Program provides support for SNA traffic across the channel and is analogous to the SNA support on the CMCC. The Multiprotocol Extensions option provides TCP/IP Offload plus all of the features of the IBM Route Expander software. Route Expander provides basic routing and WAN connectivity.

The IBM 3172 platform supports both ESCON and parallel channels. It offers Token Ring, Ethernet, and FDDI LAN connectivity and a variety of different WAN interfaces. However, the platform only supports five slots for connectivity.

The IBM 3172 was a capable LAN-to-mainframe connectivity device in its time. However, the hardware has not kept pace with the evolution in processors, memory, or bus architecture. The software options provided on top of OS/2 (a defunct operating system) did not provide the overall throughput and stability required in today's enterprise environments.

The CMCC is a very effective and complete replacement for the IBM 3172. The connectivity options offered on a Cisco channel-attached router far exceed those offered on the IBM 3172. The Cisco IOS Software offers many more advanced capabilities—DLSw+, SNASw, and QoS, to name just a few. A channel-attached Cisco 7000 Series router with CMCC can provides many times the throughput of a single IBM 3172. Finally, the CMCC supports all of the capabilities of the IBM 3172 (plus the TN3270 Server) and can be introduced with few changes to the host definitions.

IBM OSA

The OSA is an integrated communications adapter for S/390, ESCON-based mainframes that provides for direct attachment to Ethernet, Fast Ethernet, Token Ring, FDDI, and ATM networks. IBM previously released two versions of the OSA: OSA1, a two-card set that included two Intel 486 processors running channel offload software; and OSA2, a single-card replacement for the OSA1. The OSA2 eliminated the channel offload software offered on the OSA1.

The OSA-Express is the third iteration of OSA. With the availability of IBM's high-speed OSA-Express, customers are questioning when it is appropriate to use the OSA-Express, and when to use the CMCC.

The following sections describe the OSA-Express, the problems that the OSA-Express solves, and the problems it does not solve.

What Is the OSA-Express?

The OSA-Express is a network interface card (NIC) for the mainframe. It provides Gigabit Ethernet, Fast Ethernet, Ethernet, and ATM access to the LAN. The OSA-Express is important for many reasons:

- It takes advantage of new hardware architectures and operating system features to provide high-speed TCP/IP access to the mainframe.
- It removes most of the bottlenecks associated with older channel protocols and allows access into an OS/390 that is similar to access in any open systems UNIX or Windows workstation complex.
- For TCP/IP access, it removes the requirement for a channel-attached controller, such as the FEP or the CMCC.

Compared to the older OSA cards, the OSA-Express utilizes some of the following technologies:

- *Self Timed Interface (STI) bus*—The STI bus is the main data bus for the mainframe CPUs. It operates at 333 MBps. The OSA cards are directly attached to this bus in a chained fashion so there are two paths between any OSA card and the processors.
- *Queued Direct Input Output (QDIO) subsystem*—QDIO allows the OSA-Express to communicate on the STIs without the limitations of the ESCON or bus and tag channel protocols by removing the 17-MB limitation of the ESCON channel architecture.
- *Direct Memory Access (DMA) Protocol*—The CS/390 TCP/IP stack accesses the data in the OSA-Express buffers directly, without multiple data copies.
- *IP Assist features*—The OSA-Express offloads some of the processing from the mainframe TCP/IP stack, including MAC handling, packet filtering, IP multicast, and maintenance of the IP address table.
- *Configuration through the mainframe TCP/IP stack*—When you use the OSA-Express only for TCP/IP, the OSA-Express is automatically configured from information provided by the mainframe TCP/IP stack. This feature removes one of the limitations with the older OSA cards. You must still use the OSA Support Facility (OSA/SF) to configure SNA protocols and ATM cards.
- *Manageability*—The newest versions of OS/390 provide RMF and SNMP management of the OSA-Express.

When to Use the OSA-Express

The OSA-Express is becoming the method of choice for connecting the S/390 to a TCP/IP network. The OSA-Express architecture removes the limitations of the channel protocols and places the S/390 on the same plane as the large UNIX servers. By using QDIO to the STI bus, the Gigabit Ethernet and Fast Ethernet cards have direct access to the 333-MBps CPU buses. The OSA-Express is considerably faster than the current ESCON technology.

By rewriting the TCP/IP stack to use the DMA protocol against the OSA buffers, IBM has eliminated many of the buffer copies, which results in better throughput and reduced CPU resource consumption. IBM reduced the amount of configuration that is required for TCP/IP pass-through by loading the parameters from the TCP/IP profiles dataset, which eliminates the need to use the OS/2 or Windows-based OSA/SF facility.

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

In general, use the OSA-Express for high-speed TCP/IP access and use a Cisco network to gain the greatest advantages from the OSA-Express.

However, not all customers will benefit from the OSA-Express. To benefit from the OSA-Express, your system must meet the following requirements:

- *IBM mainframe*—You must be using an IBM mainframe. The OSA-Express is not supported on non-IBM hardware.
- *Mainframe must be Generation 5 or later*—The OSA-Express is not supported on IBM processors prior to Generation 5. The Generation 5 mainframes became available in 1998.
- *OS/390 Version 2, Release 7 or later*—The operating system must be Version 2, Release 7 or later. QDIO, which enables the OSA-Express performance, was introduced in Version 2, Release 7.
- *IBM CS/390 TCP/IP stack*—You must be using the IBM TCP/IP stack.
- *TCP/IP access*—The QDIO and DMA improvements in the OSA-Express are for TCP/IP traffic only. SNA traffic is supported as long as it is through SNASw or APPN/HPR/IP. Legacy SNA traffic support is provided on the lower-speed OSA-Express cards, but it does not use QDIO and DMA and represents an expensive use of a valuable ESCON card cage slot.

Use the OSA Express as a high-speed channel into the mainframe. OSA-Express is not a router or a switch, and you must connect the OSA-Express to the network through a router or a switch. OSA-Express is a good choice for accessing the TN3270 Server application on the mainframe and for high-speed FTP traffic and for HPR/IP traffic.

When Not to Use the OSA-Express

In addition to supporting high-speed TCP/IP access through Gigabit Ethernet and Fast Ethernet interfaces, the OSA-Express supports lower-speed LAN interfaces, such as 4-Mbps Token Ring, 10-Mbps Ethernet, Fast Ethernet, and ATM. OSA-Express supports native SNA traffic in the same manner as the OSA2 card. When used in the OSA-Express, the lower-speed interface cards run at wire speed, which is not true of these same interface cards in the OSA2.

Even at wire-speed, OSA cards do not represent a good use of valuable mainframe real estate for SNA traffic. The OSA cards are placed in a slot in the card cage on the mainframe. This card cage has a limited number of slots that you can use to connect a limited number of devices. You can use a slot in the card cage for either ESCON or OSA connectivity. A single slot can support four ESCON ports. The maximum theoretical throughput of these four ports is 68 MBps (4 x 17 MBps).

If you use this same slot for a Gigabit Ethernet OSA-Express, the maximum theoretical throughput is close to 120 MB, which is a good tradeoff. If this same slot were used for a 10-Mbps Ethernet OSA card, the throughput would be less than 1.5 MBps, which is not a good tradeoff. The same reasoning applies to the Fast Ethernet versions of the OSA-Express.

In these situations, you should use a router to aggregate WAN connections and to send aggregated data through a CIP or CPA.

To enable SNA support, you must use OSA/SF, an OS/2 and Microsoft Windows-based graphical configuration tool. (OSA/SF also has a 3270-style REXX interface.) The chief complaint of customers is that using the OSA/SF is cumbersome.

When to Use a CMCC

You should use a CMCC in the following situations:

- *Non-IBM mainframes*—Use the CMCC with any mainframe that supports the ESCON or bus and tag channel protocols, which is 100 percent of all IBM and PCM boxes. The OSA-Express is not an option for non-IBM mainframes.
- *Older mainframes prior to Generation 5*—Use the CMCC on the approximately 60 percent of mainframes that do not support the OSA-Express.

- *Older operating system releases prior to Version 2, Release 7*—Use the CMCC with any currently supported operating system release.
- *Aggregation of TCP/IP and SNA traffic*—The CMCC is an efficient use of the I/O card cage resources. Use the interface cards in the router to aggregate WAN traffic and to efficiently transport the combined traffic across the ESCON or bus and tag channel.
- *Offload processing*—Use the dedicated CPU and memory of the CMCC to offload processing from both the router and the mainframe. You can use the TN3270 Server application to offload the protocol conversion duties from the mainframe. You can also use the TCP/IP Offload function to offset the inefficiencies associated with the mainframe TCP/IP stack in older releases (Version 2, Release 4 and earlier).

Why You Should Use the OSA-Express with a Cisco Network

You should use the OSA-Express when your environment allows you to do so. The OSA-Express provides high-speed TCP/IP access to the mainframe; however, it functions only as a network interface. The OSA-Express does not provide routing or switching, so it is dependent on the external network to perform routing and switching.

This section describes the advantages of using a Cisco network in front of the OSA-Express. A Cisco network adds value to the OSA-Express in the following four major areas:

- Jumbo frame support
- End-to-end QoS
- Load balancing through MNLB features
- Redundancy

Jumbo Frame Support

A Jumbo Ethernet frame is 8992 bytes and is typically referred to as a 9-KB frame. A normal Ethernet frame is 1492 bytes and is typically referred to as 1500 bytes. Using the Jumbo Ethernet frame can result in significant increases in throughput (up to 75 percent) for bulk data transfers, such as file transfers or storage backups. To use Jumbo frames, all network devices between the servers must support the larger frame size.

One benefit of using a Cisco switched network with the OSA-Express is that the Gigabit Ethernet interfaces on the Catalyst 6500 switches support Jumbo Ethernet frames. The Gigabit Ethernet interface card supported by the OSA-Express also supports Gigabit Ethernet interface cards on other IBM servers, such as the RS/6000, which also supports Jumbo Ethernet frames.

End-to-End QoS

The Service Policy System of the OS/390 operating system allows you to prioritize application traffic by application name, time of day, origin, destination address pairs, and so on. The Workload Manager uses this information to prioritize dispatching applications. The OSA-Express also uses this information to prioritize outbound traffic. The OSA-Express uses one inbound traffic queue and four outbound traffic queues. The outbound traffic queues are associated with the four TCP/IP ToS precedence settings.

In conjunction with the Service Policy System, the S/390 can prioritize application performance, as well as outbound traffic.

Cisco and IBM have conducted interoperability tests. These tests prove that a Cisco network recognizes the prioritization of traffic from the OSA-Express and enforces this priority end to end through the network. The Cisco network uses features such as WFQ, Custom Queuing, Priority Queuing, and WRED to enforce the prioritization of the application traffic.

Note: In Cisco routers, WFQ is on by default on serial interfaces at speeds up to T1/E1 rates, which means you do not need to do anything extra to support the end-to-end prioritization of application traffic.

Load Balancing

Cisco and IBM have conducted testing that proves that a Cisco network can make load-balancing decisions based on information from the IBM Workload Manager. Using the MNLB feature of Cisco LocalDirector, the Cisco network can direct traffic to one or more OSA-Express adapters based upon the capability of the connected mainframe LPARs to conduct work.

Testing was conducted using a Cisco workload agent, which communicates with the IBM Workload Manager to set a metric value for each connected interface. The Cisco LocalDirector uses this information to route traffic to the connected OSA-Express adapters or CMCC cards.

Redundancy

One of the greatest differentiators for the IBM mainframe environment is its unrivalled availability. IBM mainframe environments are known for 99.99 percent availability, which equates to less than 15 minutes of downtime per year. If you are going to have an application server with this type of uptime characteristic, you want to attach it to a network with the same reliability.

Testing conducted by IBM and Cisco has shown that you can design networks to eliminate any single point of failure in the network and in the application server complex. During testing, which used both CMCCs and OSA cards, the following scenarios were created and the application availability results noted:

- *CPU failure*—Application traffic was rerouted to another server in the complex
- *LPAR failure*—Application traffic was rerouted to another server in the complex
- *Router failure*—Application traffic was rerouted through another router in the network
- *OSA failure*—Application traffic was rerouted through another OSA to the application VIPA address
- *CMCC failure*—Application traffic was rerouted through another CMCC to the application VIPA address

Summary of CMCC and IBM Channel Controllers

There are three main IBM solutions that attach the mainframe to the enterprise network—the 3745/3746 FEP, the 3172 Interconnect Controller, and the OSA. These products span generations of networking technology.

The IBM 3745/3746 FEP is the oldest, most mature IBM mainframe channel connectivity device. Groomed and enhanced over a period of more than a decade, the FEP has a great deal of functionality for SNA networks. Some of the features it provides cannot be provided by a CMCC or any other channel device—SNI, for example. However, the majority of SNA and TCP/IP traffic can very effectively and easily be supported on the CMCC. Enterprises should evaluate their needs carefully. The goal should be to keep and maintain only the smallest number of FEPs required to do FEP-specific functionality. The rest of the SNA and TCP/IP traffic will be better served by gaining access to the mainframe via a CMCC.

The IBM 3172 Interconnect Controller was a very effective LAN-to-mainframe connectivity device in the late 1980s and early 1990s. However, it has failed to keep up with technology and is no longer even sold by IBM. In many environments, the IBM 3172 cannot support the level of traffic required in enterprises today. The CMCC is a complete, cost-effective, and high-throughput replacement this device.

The OSA-Express is the latest generation of the NICs for the mainframe. The OSA-Express provides up to Gigabit Ethernet TCP/IP access to the mainframe. When you have the option to use this technology, you should do so. You also should use a Cisco network to interface to the OSA-Express card to the advantage of Jumbo frame support, end-to-end QoS, load balancing, and redundancy.

Many situations exist in which you cannot or should not use the OSA-Express. These situations include the use of non-IBM mainframes, older IBM mainframes, older operating system releases, and the need to aggregate SNA and TCP/IP traffic. In these instances, the CMCC remains the most beneficial means of providing high-speed access into the mainframe.