

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.

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.

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

Data Center Availability

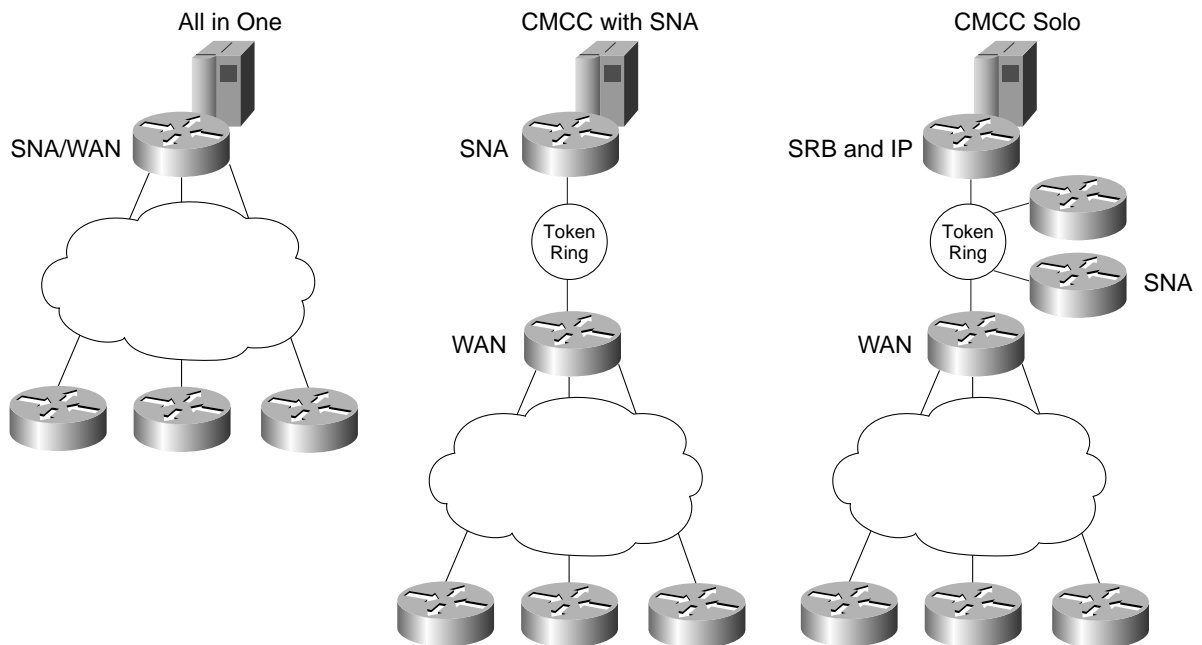
Running multiple functions and CMCCs in a single router can affect availability. If you place all your functionality (CMCCs, DLSw+, APPN/SNASw, Frame Relay traffic shaping, and so on) in a single router, you increase the likelihood of an impact due to a planned or unplanned outage in that router.

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

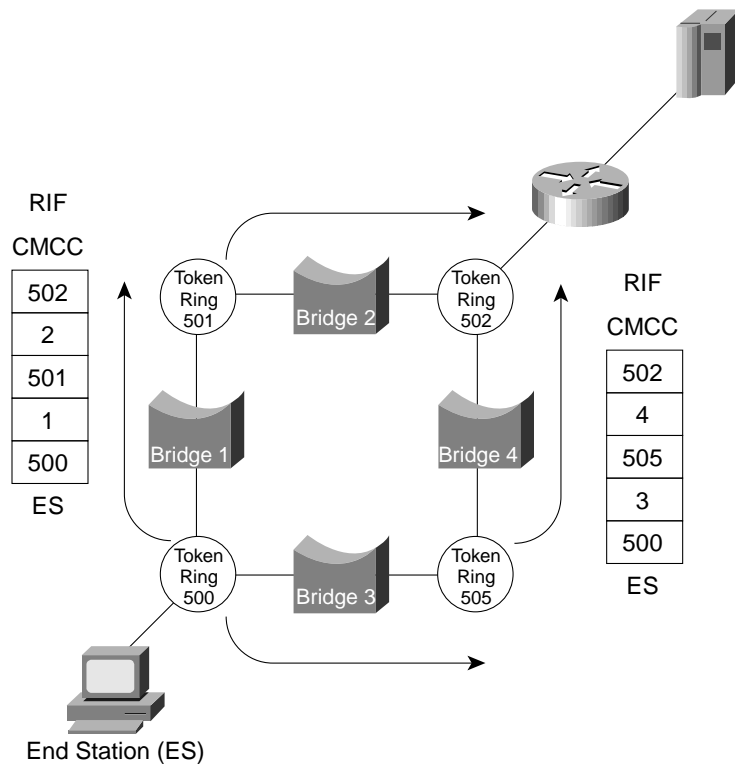
Because the CMCC appears as a LAN port and the attaching SNA end system appears as part of a switched major node, the CMCC takes advantage of the redundancy features inherent in LANs and SRB. Cisco DLSw+ provides SRB availability characteristics for devices on Ethernet or SDLC.

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.

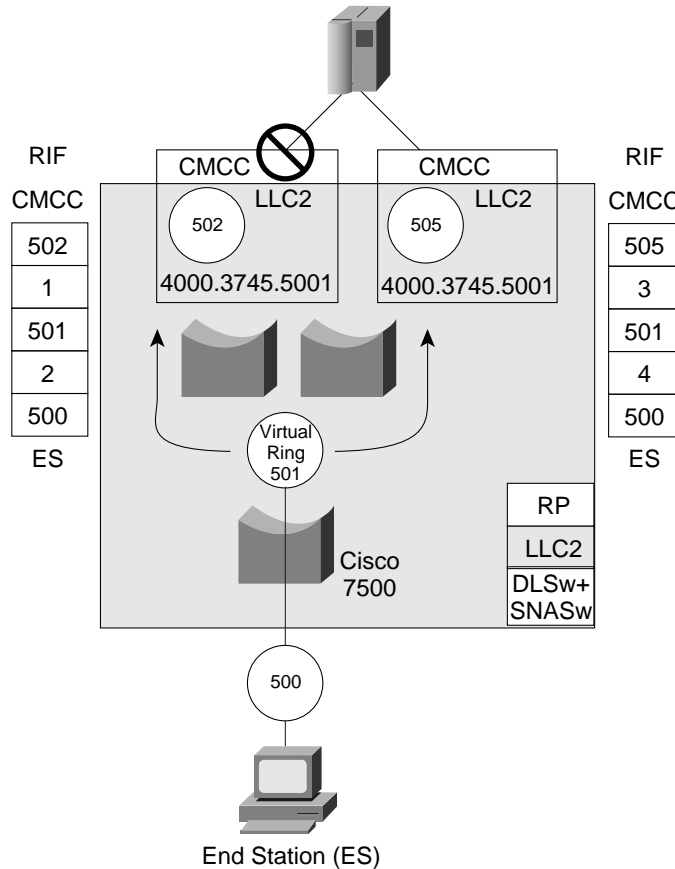
The end station then sends an XID to the CMCC using this SRB path. The CMCC forwards the XID to VTAM. The XID contains the IDBLK and IDNUM and must match an entry defined in VTAM for the session to be established.

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



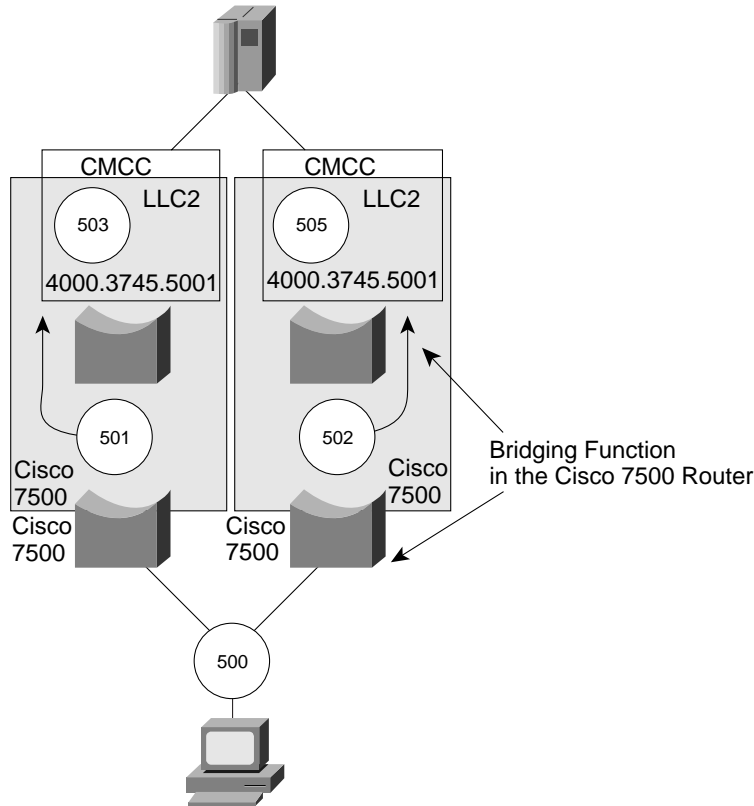
Duplicate addresses are not allowed on a single LAN segment. Duplicate addresses on different segments can be concurrently active. Note that in the case of the CMCC, these segments (as indicated by rings 502 and 505 in Figure 3-3) can be logical ring segments. (In the case of FEPs, these would be physical 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.

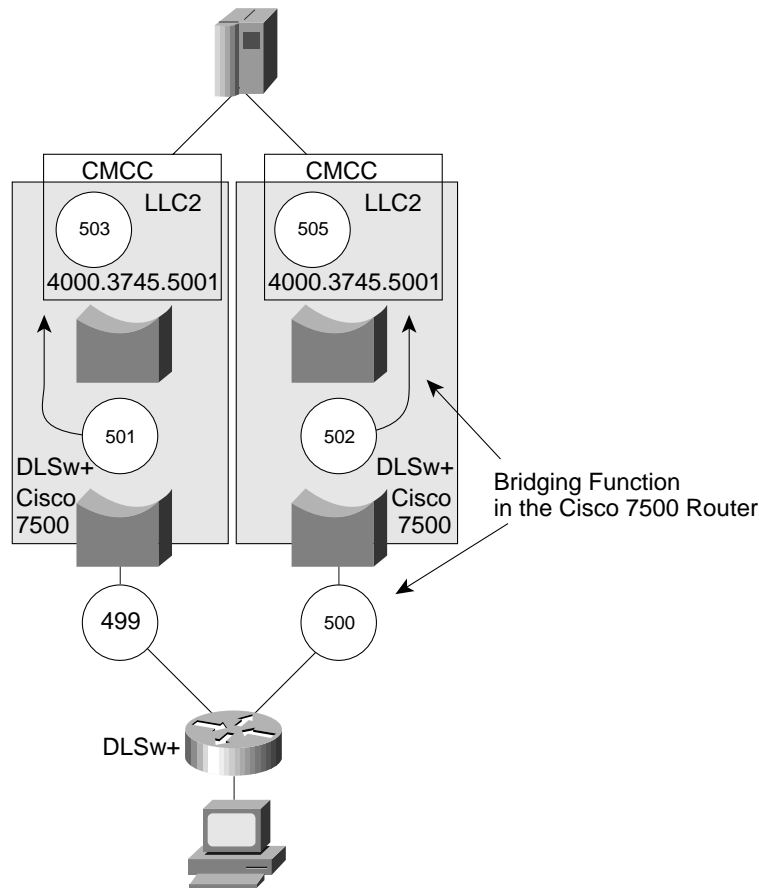
Figure 3-4 Load Balancing Using Duplicate MAC Addresses and SRB



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 **cmcp**. When using the **csna** command, the CMCC appears to VTAM as an XCA. When using the **cmcp** 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.

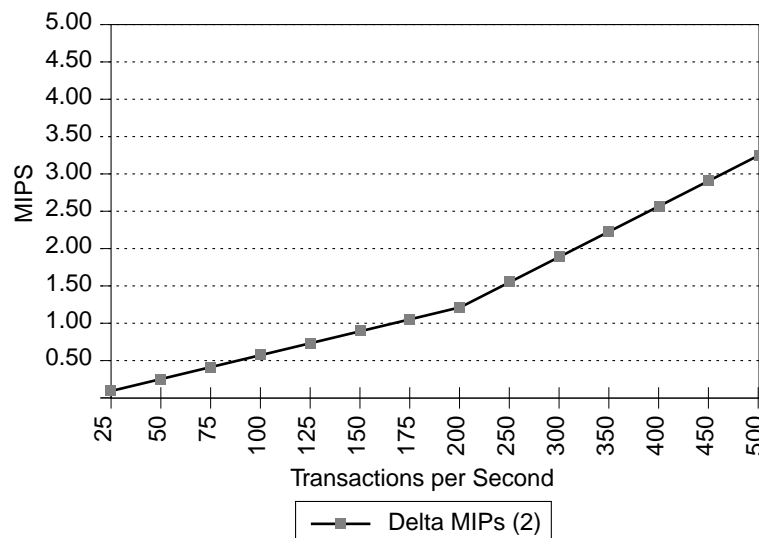
In addition, the delta MIPS required to handle a given transaction load is smaller if the transaction rate is higher, partly due to coat-tailing at higher loads, and because the VTAM service request block performs more work under each dispatch (that is, it is able to process more PIUs before releasing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