

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

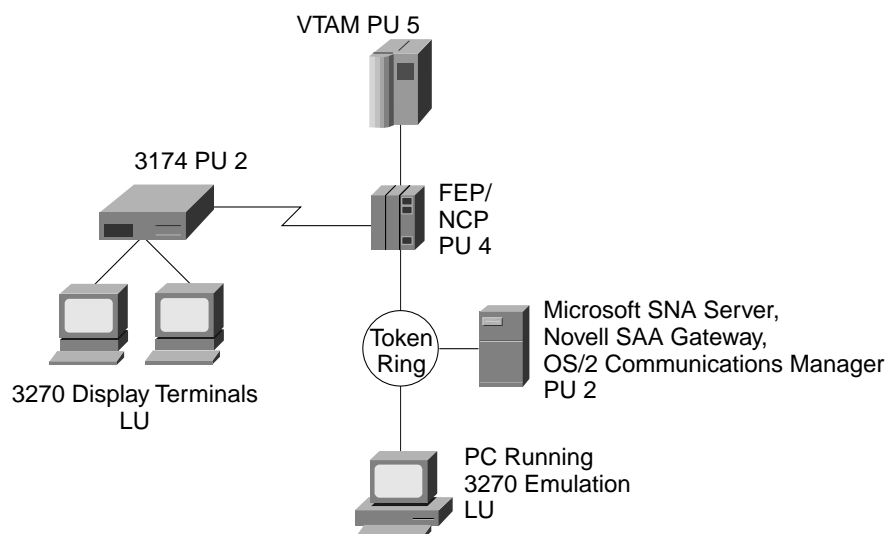
In general, all devices within a VTAM’s domain of control had to be configured to that VTAM. Later VTAM releases allowed dynamic configuration of resources using generic definitions. A VTAM can also dynamically find resources owned by another VTAM domain; such resources are known as cross-domain resources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

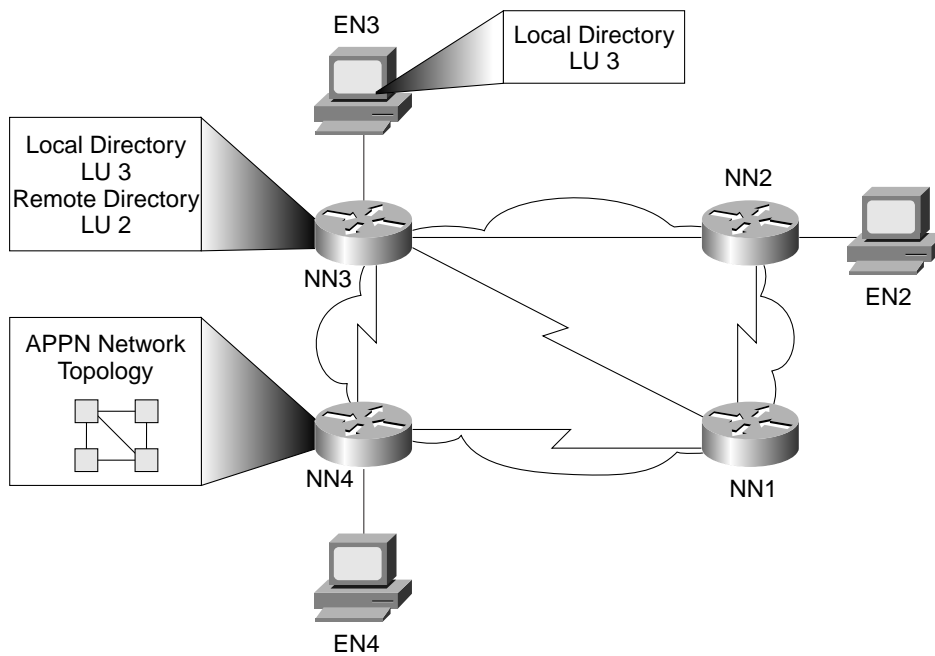
- *Network nodes (NNs)*—These nodes are SNA routers, responsible for locating resources, selecting paths, and working with the users to set up sessions.
- *End nodes (ENs)*—These nodes are application hosts, end users, or controllers representing multiple users.
- *Low-entry networking (LEN) nodes*—These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



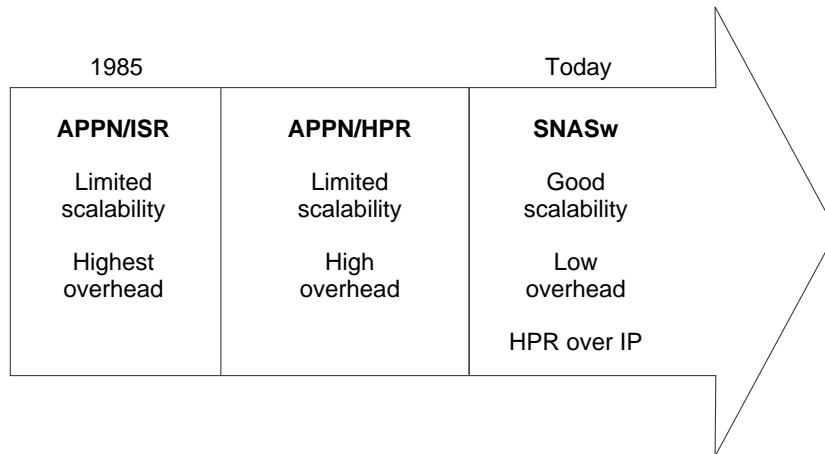
APPN ENs provide local directory services and communicate with their NN server to access other resources in the network. APPN ENs dynamically register their local resources with their upstream NN server.

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



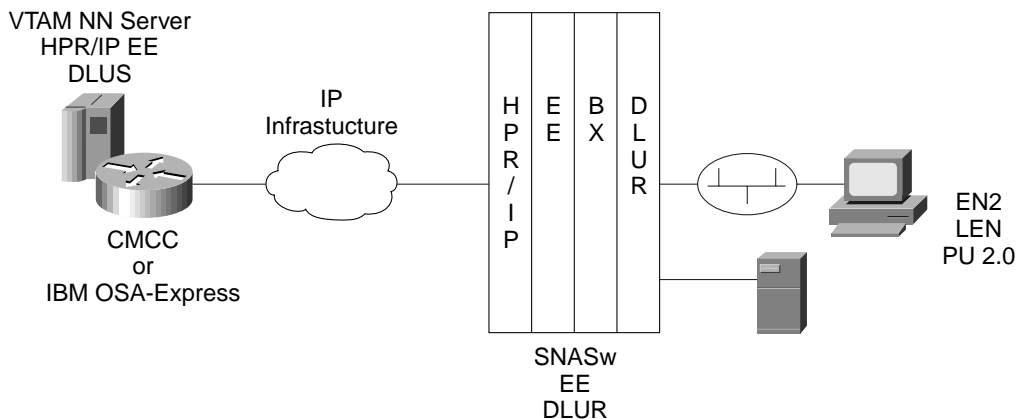
Overview of the SNASw DLUR/DLUS Features

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.

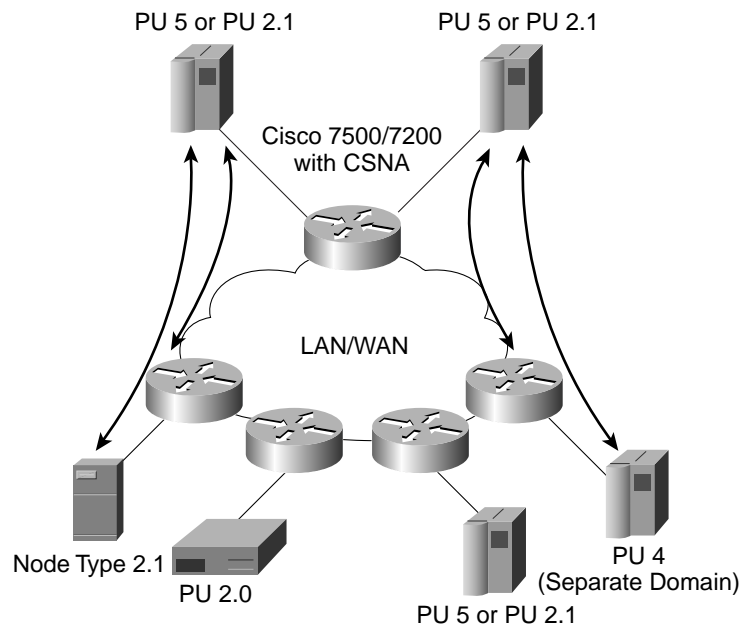
Figure 2-6 SNASw DLUR Feature



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.)

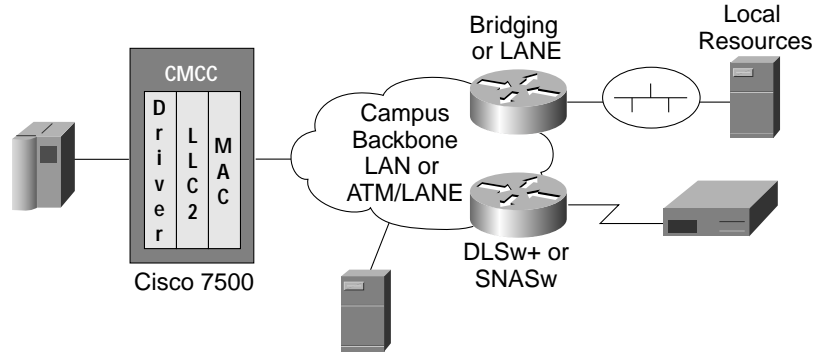
Figure 2-7 Connectivity among SNA Devices



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.

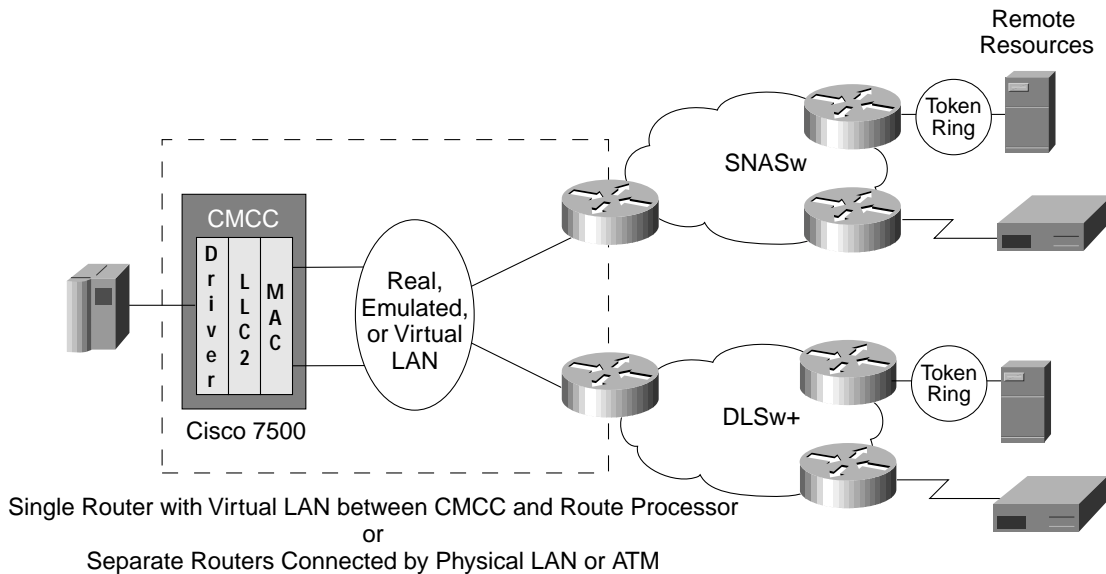
Local LAN-attached resources or campus resources connected to an Asynchronous Transfer Mode (ATM) LAN Emulation (LANE) backbone can bridge into the CMCC-attached router. Local Synchronous Data Link Control (SDLC) devices can attach directly to a Cisco router and use either DLSw+ local switching or SNASw 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



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

