

Designing APPN Internetworks

Advanced peer-to-peer networking (APPN) is a second generation of the Systems Network Architecture (SNA) from IBM. It moves SNA from a hierarchical, mainframe-centric environment to a peer-to-peer environment. It provides capabilities similar to other LAN protocols, such as dynamic resource definition and route discovery.

This chapter focuses on developing the network design and planning a successful migration to APPN. It covers the following topics:

- Evolution of SNA
- When to Use APPN as Part of a Network Design
- When to Use APPN Versus Alternative Methods of SNA Transport
- Overview of APPN
- Scalability Issues
- Backup Techniques in an APPN Network
- APPN in a Multiprotocol Environment
- Network Management
- Configuration Examples

Note Although this chapter does discuss using APPN with DLSw+, for detailed information on using DLSw+, refer to Chapter 7, “Designing DLSw+ Internetworks.”

Evolution of SNA

Introduced in 1974, subarea SNA made the mainframe computer running Advanced Communications Function/Virtual Telecommunication Access Method (ACF/VTAM) the hub of the network. The mainframe was responsible for establishing all sessions (a connection between two resources over which data can be sent), activating resources, and deactivating resources. The design point of subarea SNA was reliable delivery of information across low-speed analog lines. Resources were explicitly predefined. This eliminated the need for broadcast traffic and minimized header overhead.

Many enterprises today maintain two networks: a traditional, hierarchical SNA subarea network and an interconnected LAN network that is based on connectionless, dynamic protocols. The advantage of the subarea SNA network is that it is manageable and provides predictable response time. The disadvantages are that it requires extensive system definition and does not take advantage of the capabilities of intelligent devices (for example, the PCs and workstations).

Role of APPN

With APPN, you can consolidate the two networks (an SNA subarea network and an interconnected LAN network) because APPN has many of the characteristics of the LAN networks and still offers the advantages of an SNA network. The major benefits of using APPN include the following:


- Connections are peer-to-peer, allowing any end user to initiate a connection with any other end user without the mainframe (VTAM) involvement.
- APPN supports subarea applications as well as newer peer-to-peer applications over a single network.
- APPN provides an effective routing protocol to allow SNA traffic to flow natively and concurrently with other protocols in a single network.
- Traditional SNA class of service (COS)/transmission priority can be maintained.

As SNA has evolved, one feature has remained critical to many users: COS. This feature provides traffic prioritization on an SNA session basis on the backbone. This, in turn, allows a single user to have sessions with multiple applications, each with a different COS. In APPN, this feature offers more granularity and extends this capability all the way to the end node rather than just between communication controllers.

Types of APPN Nodes

An APPN network has three types of nodes: LEN nodes, end nodes (EN), and network nodes (NN), as shown in Figure 6-1.

Figure 6-1 Different types of APPN nodes.



Note Throughout the rest of this chapter, the abbreviations EN and NN are used in the illustrations. The full terms (end node and network node) are used within the text for clarity.

Table 6-1 describes these different types of APPN nodes. The control point (CP), which is responsible for managing a node's resources and adjacent node communication in APPN, is key to an APPN node. The APPN Control Point is the APPN equivalent of the SSCP.

Table 6-1 Different Types of APPN Nodes

Type of APPN Node	Description
Local Entry Networking (LEN) nodes	LEN nodes are pre-APPN, peer-to-peer nodes. They can participate in an APPN network by using the services provided by an adjacent network node. The CP of the LEN node manages the local resources but does not establish a CP-CP session with the adjacent network node. Session partners must be predefined to the LEN node, and the LEN node must be predefined to the adjacent network node. LEN nodes are also referred to as SNA node type 2.1, physical unit (PU) type 2.1, or PU2.1.
End nodes	End nodes contain a subset of full APPN functionality. They access the network through an adjacent network node and use the adjacent network node's routing services. An end node establishes a CP-CP session with an adjacent network node, and then uses that session to register resources, request directory services, and request routing information.
Network nodes	Network nodes contain full APPN functionality. The CP in a network node is responsible for managing the resources of the network node along with the attached end nodes and LEN nodes. The CP establishes CP-CP sessions with adjacent end nodes and network nodes. It also maintains network topology and directory databases, which are created and updated by dynamically gathering information from adjacent network nodes and end nodes over CP-CP sessions. In an APPN environment, network nodes are connected by transmission groups (TGs), which in the current APPN architecture refers to a single link. Consequently, the network topology is a combination of network nodes and transmission groups.

For more background information on APPN, refer to the section "Overview of APPN" later in this chapter.

When to Use APPN as Part of a Network Design


APPN has two key advantages over other protocols:

- Native SNA routing
- COS for guaranteed service delivery

APPN, like Transmission Control Protocol/Internet Protocol (TCP/IP), is a routable protocol in which routing decisions are made at the network nodes. Although only the network node adjacent to the originator of the session selects the session path, every network node contributes to the process by keeping every other network node informed about the network topology. The network node adjacent to the destination also participates by providing detailed information about the destination. Only routers that are running as APPN network nodes can make routing decisions.

You need APPN in your network when a routing decision (for example, which data center or path) must be made. Figure 6-2 helps to illustrate the criteria you use to determine where APPN should be used in a network.

Figure 6-2 Determining where to use APPN in a network.



In Figure 6-2, a single link connects the branch office to the backbone. Therefore, a routing decision does not need to be made at the branch office. Consequently, an APPN network node might not be necessary at those sites.

Because there are two data centers, however, the routing decision about which data center to send the message to must be made. This routing decision can be made either at the data center or at the backbone routers. If you want this routing decision made at the data center, all messages are sent to a single data center using DLSw+, for example, and then routed to the correct data center using APPN only in the routers in the data center. If you want the routing decision to be made at the backbone routers, place the APPN network node in the backbone routers, where alternative paths are available for routing decisions outside of the data center. In this example, this latter approach is preferred because it isolates the function at the data centers routers to channel attachment, reduces the number of hops to the second data center, and provides a path to a backup data center if something catastrophic occurs.

Because APPN requires more memory and additional software, it is generally a more expensive solution. The advantages of direct APPN routing and COS, however, often offset the added expense. In this case, the added expense to add APPN to the backbone and data center routers might be justifiable, whereas added expense at the branch might not be justifiable.

APPN at Every Branch

There are two cases for which adding an APPN network node at every branch can be cost justified:

- When COS Is Required
- When Branch-to-Branch Routing Is Required

When COS Is Required

COS implies that the user accesses multiple applications and must be able to prioritize traffic at an application level. Although other priority schemes, such as custom queuing, might be able to prioritize by end user, they cannot prioritize between applications for an individual user. If this capability is critical, APPN network nodes must be placed in the individual branches to consolidate the traffic between multiple users using COS. For instance, COS can ensure that credit card verification always gets priority over batch receipts to a retail company's central site.


It is important to understand where COS is used in the network today. If the network is a subarea SNA network, COS is used only between front-end processors (FEPs) and ACF/VTAM on the mainframe. Unless there is already an FEP at the branch office, they do not have traffic prioritization from the branch, although traffic can be prioritized from the FEP out. In this case, adding an APPN network node at the branch office would prioritize the traffic destined for the data center sooner rather than waiting until it reaches the FEP—adding function over what is available today.

When Branch-to-Branch Routing Is Required

If branch-to-branch traffic is required, you can send all traffic to the central site and let those APPN network nodes route to the appropriate branch office. This is the obvious solution when both data center and branch-to-branch traffic are required and the branch is connected to the backbone over a single link. However, if a separate direct link to another branch is cost-justifiable, routing all traffic to the data center is unacceptable. In this case, making the routing decision at the branch is necessary. Using an APPN network node at the branch, data center traffic is sent over the data center link and branch-to-branch traffic is sent over the direct link.

In the example in Figure 6-3, each branch has two links to alternative routers at the data center. This is a case where APPN network nodes might be required at the branches so that the appropriate link can be selected. This can also be the design for branch-to-branch routing, adding a single hop rather than creating a full mesh of lines. This provides more direct routing than sending everything through the data center.

Figure 6-3 Sample network for which branch-to-branch routing is required.



As you also learn in this chapter, scalability issues make it advantageous to keep the number of network nodes as small as possible. Understanding where native routing and COS is needed is key in minimizing the number of network nodes.

In summary, choosing where to implement APPN must be decided based on cost, scalability, and where native routing and COS are needed. Implementing APPN everywhere in your network might seem to be an obvious solution, even when not necessary. It must be understood, however, that if you were to deploy APPN everywhere in your network it probably would be a more costly solution than necessary and may potentially lead to scalability problems. Consequently, the best solution is to deploy APPN only where it is truly needed in your network.

When to Use APPN Versus Alternative Methods of SNA Transport

APPN and boundary network node (BNN)/boundary access node (BAN) over Frame Relay using RFC 1490 are the two methods of native SNA transport, where SNA is not encapsulated in another protocol. BAN and BNN allow direct connection to an FEP, using the Frame Relay network to switch messages, rather than providing direct SNA routing.

Although *native* might seem to be the appropriate strategy, APPN comes at the price of cost and network scalability, as indicated in the preceding section. With BNN/BAN additional cost is required to provide multiprotocol networking because the FEP does not handle multiple protocols. This implies that additional routers are required in the data center for other protocols and separate virtual circuits are required to guarantee service delivery for the SNA or APPN traffic.

DLSw+ provides encapsulation of SNA, where the entire APPN message is carried as data inside a TCP/IP message. There is often concern about the extra 40 bytes of header associated with TCP/IP. However, because Cisco offers alternatives such as Data Link Switching Lite, Fast Sequenced Transport (FST), and Direct Transport, which have shorter headers, header length is deemed noncritical to this discussion.

DLSw+ is attractive for those networks in which the end stations and data center will remain SNA-centric, but the backbone will be TCP/IP. This allows a single protocol across the backbone, while maintaining access to all SNA applications. DLSw+ does not provide native APPN routing, nor does it provide native COS. Consequently, DLSw+ is preferable for networks, in which cost is a key criterion, that have the following characteristics:

- A single data center or mainframe
- Single links from the branches

In general, DLSw+ is a lower-cost solution that requires less memory and software. In the vast majority of networks, DLSw+ will be combined with APPN—using APPN only where routing decisions are critical. With TCP/IP encapsulation, the TCP layer provides the same reliable delivery as SNA/APPN, but does not provide the native routing and COS.

TN3270 transports 3270 data stream inside a TCP/IP packet without SNA headers. Therefore, this solution assumes that the end station has only a TCP/IP protocol stack and no SNA. Therefore, TN3270 is not an alternative to APPN because APPN assumes the end station has an SNA protocol stack. APPN, like DLSw+, may still be required in the network to route between TN3270 servers and multiple mainframes or data centers.

In summary, APPN will frequently be used with DLSw+ in networks when a single backbone protocol is desired. BAN/BNN provides direct connectivity to the FEP but lacks the multiprotocol capabilities of other solutions. TN3270 is used only for TCP/IP end stations.

Overview of APPN

This section provides an overview of APPN and covers the following topics:

- Defining Nodes
- Establishing APPN Sessions
- Understanding Intermediate Session Routing
- Using Dependent Logical Unit Requester/Server

Defining Nodes

Nodes, such as ACF/VTAM, OS/400 and Communications Server/2 (CS/2), can be defined as either network nodes or end nodes. When you have a choice, consider the following issues:

- *Network size*—How large is the network? Building large APPN networks can introduce scalability issues. Reducing the number of network nodes is one solution for avoiding scalability problems. For more information on reducing the number of network nodes, see the section “Reducing the Number of Network Nodes” later in this chapter.
- *Role of the node*—Is it preferable to have this node performing routing functions as well as application processing? A separate network node can reduce processing cycles and memory requirements in an application processor.

Generally, you should define a network node whenever a routing decision needs to be made.

APPN Node Identifiers

An APPN node is identified by its network-qualified CP name, which has the format netid.name. The network identifier (netid) is an eight-character name that identifies the network or subnetwork in which the resource is located. The network identifier and name must be a combination of uppercase letters (A through Z), digits (0 through 9), and special characters (\$,#,or @) but cannot have a digit as the first character.

Establishing APPN Sessions

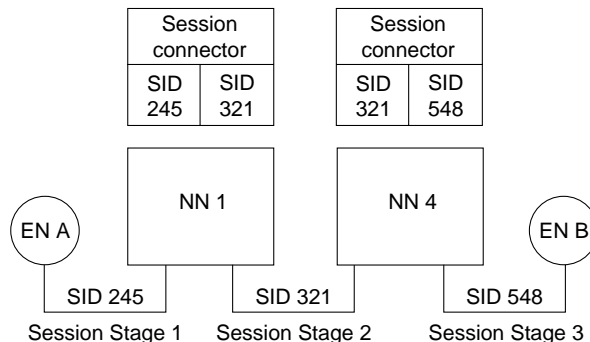
In order for an APPN session to be established, the following must occur:

- 1 The end user requests a session with an application, which causes the end node to begin the process of session establishment by sending a LOCATE message to its network node server. For session initiation, the network node server provides the path to the destination end node, which allows the originating end node to send messages directly to the destination.
- 2 The network node uses directory services to locate the destination by first checking its internal directories. If the destination is not included in the internal directory, the network node sends a LOCATE request to the central directory server if one is available. If a central directory server is not available, the network node sends a LOCATE broadcast to the adjacent network nodes that in turn propagate the LOCATE throughout the network. The network node server of the destination returns a reply that indicates the location of the destination.
- 3 Based on the location of the destination, the COS requested by the originator of the session, the topology database, and the COS tables, the network node server of the originator selects the least expensive path that provides the appropriate level of service.
- 4 The originating network node server sends a LOCATE reply to the originating end node. The LOCATE reply provides the path to the destination.
- 5 The originating end node is then responsible for initiating the session. A BIND is sent from the originating end node to the destination end node, requesting a session. After the destination replies to the BIND, session traffic can flow.

Understanding Intermediate Session Routing

Session connectors are used in place of routing tables in APPN. The unique session identifier and port from one side of the node are mapped to the unique session identifier and port on the other side. As data traffic passes through the node, the unique session identifier in the header is swapped for the outgoing identifier and sent out on the appropriate port, as shown in Figure 6-4.

Figure 6-4 Intermediate session routing label swap.



This routing algorithm is called *intermediate session routing* (ISR). It supports dynamic route definition and incorporates the following legacy features:

- *Node-to-node error and flow control processing*—This reflects the 1970s method of packet switching in which many line errors dictated error and flow control at each node. Given the current high-quality digital facilities in many locations, this redundant processing is unnecessary and significantly reduces end-to-end throughput. End-to-end processing provides better performance and still delivers the necessary reliability.
- *Disruptive session switching around network failures*—Whenever a network outage occurs, all sessions using the path fail and have to be restarted to use an alternative path.

Because these features are undesirable in most high-speed networks today, a newer routing algorithm—High Performance Routing (HPR)—has been added to APPN that supports nondisruptive rerouting around failures and end-to-end error control, flow control, and segmentation. Cisco routers support both ISR and HPR.


Using Dependent Logical Unit Requester/Server

Dependent Logical Unit Requester/Server (DLUR/DLUS) is an APPN feature that allows legacy traffic to flow on an APPN network. Prior to the introduction of this feature, the APPN architecture assumed that all nodes in a network could initiate peer-to-peer traffic (for example, sending the BIND to start the session). Many legacy terminals that are referred to as Dependent Logical Units (DLUs) cannot do this and require VTAM to notify the application, which then sends the BIND.

Getting the legacy sessions initiated requires a client-server relationship between ACF/VTAM (Dependent LU server—DLUS) and the Cisco router (Dependent LU Requester—DLUR). A pair of logical unit (LU) type 6.2 sessions are established between the DLUR and DLUS—one session is established by each end point. These sessions are used to transport the legacy control messages that must flow to activate the legacy resources and initiate their logical unit to logical unit (LU-LU) sessions. An LU-LU session is the connection that is formed when the five steps described earlier in the section “Establishing APPN Sessions” are completed.

For example, an activate logical unit (ACTLU) message must be sent to the LU to activate a legacy LU. Because this message is not recognized in an APPN environment, it is carried as encapsulated data on the LU 6.2 session. DLUR then deencapsulates it, and passes it to the legacy LU. Likewise, the DLU session request is passed to the ACF/VTAM DLUS, where it is processed as legacy traffic. DLUS then sends a message to the application host, which is responsible for sending the BIND. After the legacy LU-LU session is established, the legacy data flows natively with the APPN traffic, as shown in Figure 6-5.

Figure 6-5 DLU session processing.



Cisco Implementation of APPN

This section provides an overview of Cisco's implementation of APPN and discusses where APPN resides in the Cisco IOS software. Cisco licensed the APPN source code from IBM and then ported it to the Cisco IOS software using network services from the data-link controls (DLCs).

Applications use APPN to provide network transport. APPN runs on top of the Cisco IOS software. APPN is a higher-layer protocol stack that requires network services from DLC. Cisco's APPN implementation is compliant with the APPN Architecture of record. When used with other features in the Cisco IOS software, APPN provides the following unique features:

- APPN can use DLSw+ or RSRB as a network transport, thereby supporting APPN over a native TCP/IP network.
- APPN can be used with downstream physical unit concentration (DSPU) to reduce the number of downstream PUs visible to VTAM. This reduces VTAM definition and network restart times.
- In addition to COS, priority queuing, custom queuing, and weighted fair queuing can be used with COS to ensure traffic prioritization and/or bandwidth reservation between protocols.
- Network management options are supported that include native SNA management services using Native Service Point (NSP) in the Cisco router, and Simple Network Management Protocol (SNMP) management using CiscoWorks Blue applications.
- Using Channel Interface Processor (CIP) or Channel Port Adapter (CPA), the Cisco APPN network node can interface directly with ACF/VTAM across the channel. VTAM can be defined either as an end node or network node.

Scalability Issues

As a single-network link state architecture, the network topology is updated as changes occur. This results in significant network traffic if instability occurs, and significant memory and processing to maintain the large topology databases and COS tables. Similarly, in large networks, dynamic discovery of resources can consume significant bandwidth and processing. For these reasons, scalability becomes a concern as network size increases. The number of nodes that are too large depends on the following:

- Amount of traffic
- Network stability
- The number of the techniques, which are described in this section, that are being used to control traffic and processing

Essentially, to allow growth of APPN networks, the network design must focus on reducing the number of topology database updates (TDUs) and LOCATE search requests.

Topology Database Update Reduction

APPN is a link-state protocol. Like other link-state-based algorithms, it maintains a database of the entire topology information of the network. Every APPN network node in the network sends out TDU packets that describe the current state of all its links to its adjacent network nodes. The TDU contains information that identifies the following:

- The characteristics of the sending node
- The node and link characteristics of the various resources in the network
- The sequence number of the most recent update for each described resource

A network node that receives a TDU packet propagates this information to its adjacent network nodes using a flow reduction technique. Each APPN network node maintains full knowledge of the network and how the network is interconnected. Once a network node detects a change to the network (either a change to the link, or the node), it floods TDUs throughout the network to ensure rapid convergence. If there is an unstable link in the network, it can potentially cause many TDU flows in a network.

As the number of network nodes and links increases, so does the number of TDU flows in your network. This type of distributing topology can consume significant CPU cycles, memory, and bandwidth. Maintaining routes and a large, complete topology subnet can require a significant amount of dynamic memory.

You can use the following techniques to reduce the amount of TDU flows in the network:


- Reduce the number of links
- Reduce the number of CP-CP sessions
- Reduce the number of network nodes in the network

Reducing the Number of Links

The first technique for reducing the amount of TDU flows in the network is to reduce the number of links in your network. In some configurations, it might be possible to use the concept of *connection network* to reduce the number of predefined links in your network. Because network nodes exchange information about their links, the fewer links you define, the fewer TDU flows can occur.

Figure 6-6 shows the physical view of an APPN network. In this network NN1, NN2, and NN3 are routers attached to an FDDI LAN.

Figure 6-6 Physical view of an APPN network.




The network-node server (NNS), EN1, and EN2 hosts are attached to the same FDDI LAN via a CIP router or a cluster controller. These nodes on the FDDI LAN have any-to-any connectivity. To reflect any-to-any connectivity in APPN, NN1 needs to define a link to NN2, NN3, NNS (VTAM host), EN1 (VTAM data host), and EN2 (EN data host). The transmission groups connecting network nodes are contained in the network topology database. For every link that is defined to the network node, TDUs are broadcast.

Note Throughout the rest of this chapter, the abbreviation NNS is used in the illustrations. When the text refers to an NNS icon in an illustration, the abbreviation is also used; otherwise, the full term (network-node server) is used within the text for clarity.

Figure 6-7 shows the logical view of the APPN network, shown earlier in Figure 6-6. When NN1 first joins the network, NN1 activates the links to NN2, NN3, NNS, EN1, and EN2. CP-CP sessions are established with the adjacent network nodes. Each adjacent network node sends a copy of the current topology database to NN1. Similarly, NN1 creates a TDU about itself and its links to other network nodes and sends this information over the CP-CP sessions to NN2, NN3 and NNS. When NN2 receives the TDU from NN1, it forwards the TDU to its adjacent network nodes, which are NN3 and NNS. Similarly, NN3 and NNS receive the TDU from NN1 and broadcast this TDU to their adjacent network nodes. The result is that multiple copies of the TDU are received by every network node.

Figure 6-7 Logical view of an APPN network without connection network deployed.




The transmission groups that connect the end nodes are not contained in the network topology database. Consequently, no TDUs are broadcast for the two links to EN1 and EN2. If the number of transmission groups connecting network nodes can be reduced, the number of TDU flows can also be reduced.

By using the concept of connection networks, you can eliminate the transmission group definitions, and therefore reduce TDU flows. A connection network is a single virtual routing node (VRN), which provides any-to-any connectivity for any of its attached nodes. The VRN is not a physical node, it is a logical entity that indicates that nodes are using a connection network and a direct routing path can be selected.

Figure 6-8 shows the APPN network shown in Figure 6-6 with connection network deployed.

Figure 6-8 Logical view of an APPN network with connection network deployed.



NN1, NN2, and NN3 define a link to the network-node server (NNS) and a link to the VRN. When the link between NN1 and NNS is activated, NNS sends a copy of the current network topology database to NN1. NN1 creates a TDU about itself, its link to NNS, and its link to the VRN. It then sends this information to NNS. NN1 does not have a link defined to NN2 and NN3, therefore, there are no TDUs sent to NN2 and NN3 from NN1. When NNS receives the TDU information from NN1, NNS forwards it to NN2 and NN3. Neither NN2 nor NN3 forwards the TDU information because they only have a connection to NNS. This significantly reduces the number of TDU flows in the network.

When a session is activated between resources on the connection network, the network-node server recognizes that this is a connection network and selects a direct route rather than routing through its own network nodes. Cisco recommends that you apply the concept of connection networks whenever possible. Not only does it reduce the number of TDU flows in the network, it also greatly reduces system definitions.

As shown in the example, a LAN (Ethernet, Token Ring, or FDDI) can be defined as a connection network. With ATM LAN Emulation (LANE) services, you can interconnect ATM networks with traditional LANs. From APPN's perspective, because an ATM-emulated LAN is just another LAN, connection network can be applied. In addition to LANs, the concept of connection networks can


apply to X.25, Frame Relay, and ATM networks. It should also be noted that technologies such as RSRB and DLSw appear as LANs to APPN. You can also use connection network in these environments. APPN, in conjunction with DLSw+ or RSRB, provides a synergy between routing and bridging for SNA traffic.

Reducing the Number of CP-CP Sessions

The second technique for reducing the amount of TDU flows in the network is to reduce the number of CP-CP sessions in your network. Network nodes exchange topology updates over CP-CP sessions. The number of CP-CP sessions has a direct impact on the number of TDU flows in the network.

For example, in Figure 6-9, NN2, NN3, NN4, and NN5 are in a fully meshed network. Every network node establishes CP-CP sessions with its adjacent network nodes. This means that NN2 establishes CP-CP sessions with NN3, NN4, and NN5. NN3 establishes CP-CP sessions with NN2, NN4, NN5, and so forth.

Figure 6-9 Fully meshed CP-CP sessions.



If the link fails between NN1 and NN2, TDU updates are broadcast from NN2 to NN3, NN4, and NN5. When NN3 receives the TDU update, it resends this information to NN4 and NN5. Similarly, when NN5 receives the TDU update, it resends this information to NN3 and NN4. This means that NN4 receives the same information three times. It is recommended that the number of CP-CP sessions are kept to a minimum so that duplicate TDU information will not be received.

In Figure 6-10, CP-CP sessions exist only between NN2 and NN3, NN2 and NN4, and NN2 and NN5; no other CP-CP sessions exist. When the link fails between NN1 and NN2, NN2 broadcasts transmission group updates to NN3, NN4, and NN5. None of the three NNs forwards this information to the rest of the network because CP-CP sessions do not exist. Although this minimizes the TDU flows, if the link between NN2 and NN3 fails, this becomes a disjointed APPN network and NN3 is isolated.


Figure 6-10 Single pair of CP-CP sessions.


Figure 6-11 shows a more efficient design that also provides redundancy. Every network node has CP-CP sessions with two adjacent network nodes. NN2 has CP-CP sessions with NN3 and NN5. If the link between NN2 and NN3 fails, TDU updates will be sent via NN5 and NN4.

For redundancy purposes, it is recommended that each network node has CP-CP sessions to two other network nodes if possible.

Reducing the Number of Network Nodes

The third technique for reducing the amount of TDU flows in the network is to reduce the number of network nodes by defining APPN nodes only at the edges of the network. Minimizing the number of network nodes also reduces the size of the network topology. The following are some technologies for reducing the number of network nodes:

- APPN over DLSw+
- APPN over Frame Relay Access Server (FRAS)/BNN or BAN
- APPN over RSRB

Figure 6-11 Dual pair of CP-CP sessions.


APPN over DLSw+

Data link switching is one way to reduce the number of network nodes in the network. DLSw+ is a means of transporting APPN traffic across a WAN, where APPN network nodes and/or end nodes are defined only at the edges of the network. Intermediate routing is through DLSw+ and not via native SNA.

DLSw+ defines a standard to integrate SNA/APPN and LAN internetworks by encapsulating these protocols within IP. Cisco's implementation of DLSw, known as DLSw+, is a superset of the current DLSw architecture. DLSw+ has many value-added features that are not available in other vendors' DLSw implementations. APPN, when used with DLSw, can benefit from the many scalability enhancements that are implemented in DLSw+, such as border peer, on-demand peers, caching algorithms, and explorer firewalls.

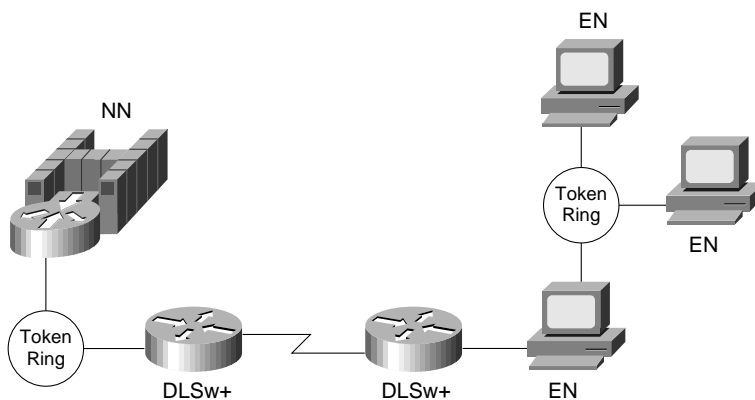
In Figure 6-12, sessions between end-node workstations and the host are transported over the DLSw+ network.

Figure 6-12 APPN with DLSw+.



VTAM acts as the network-node server for remote end-node workstations. Optionally, if multiple VTAMs or data centers exist, APPN on the channel-attached router(s) or on other routers in the data center can offload VTAM by providing the SNA routing capability, as shown in Figure 6-13.

Figure 6-13 APPN with DLSw+ using a channel-attached router.



DLSw+ also brings nondisruptive rerouting in the event of a WAN failure. Using DLSw+ as a transport reduces the number of network nodes in the network. A disadvantage is that remote end-node workstations require WAN connections for NNS services. Another disadvantage is that without APPN in the routers, APPN transmission priority is lost when traffic enters the DLSw+ network.

For detailed information on DLSw and DLSw+, refer to Chapter 7, “Designing DLSw+ Internetworks.”

APPN over FRAS BNN/BAN

If the APPN network is based on a Frame Relay network, one option is to use the FRAS/BNN or the Frame Relay BAN function for host access. Both BNN and BAN allow a Cisco router to attach directly to an FEP. When you use FRAS/BNN, you are assuming that the Frame Relay network is performing the switching and that native routing is not used within the Frame Relay network. For an example of how APPN with FRAS BNN/BAN can be used in your network design, see the section “Example of APPN with FRAS BNN” later in this chapter.

APPN over RSRB

Using RSRB, the SNA traffic can be bridged from a remote site to a data center. The use of RSRB significantly reduces the total number of network nodes in the network, thus reducing the number of TDU flows in the network. Another advantage of using RSRB is that it provides nondisruptive routing in the event of a link failure. For more information on using RSRB, refer to Chapter 4, “Designing SRB Internetworks.”

LOCATE Search Reduction

This section describes the broadcast traffic in an APPN network and how LOCATE searches can become a scalability issue in an APPN network. The impact of LOCATE searches in an APPN network varies from one network to the other. This section first identifies some of the causes of an excessive number of LOCATE searches, and then discusses the following four techniques you can use to minimize them:

- Safe-Store of Directory Cache
- Partial Directory Entries
- Central Directory Server (CDS)/Client
- Central Resource Registration

An APPN network node provides dynamic location of network resources. Every network node maintains dynamic knowledge of the resources in its own directory database. The distributed directory database contains a list of all the resources in the network. The LOCATE search request allows one network node to search the directory database of all other network nodes in the network.

When an end-node resource requests a session with a target resource that it has no knowledge of, it uses the distributed search capabilities of its network-node server to locate the target resource. If the network node does not have any knowledge of the target resource, the network node forwards the locate search request to all its adjacent network nodes requesting these nodes to assist the network-node server to locate the resource. These adjacent network nodes propagate these locate search requests to their adjacent network nodes. This search process is known as *broadcast search*.

Although several mechanisms are put into place to reduce the LOCATE broadcast searches (for example, resource registration, and resource caching), there might still be an excessive amount of LOCATE flows in a network for such reasons as the network resources no longer exist, there is a mixture of subarea networks and APPN networks, or the resources are temporarily unavailable.

Safe-Store of Directory Cache

The first technique that you can use to minimize the LOCATE flows in your APPN network is the Safe-Store of Directory Cache, which is supported by the Cisco network-node implementation. Cache entries in a network node's directory database can be periodically written to a permanent storage medium: a tftp host. This speeds recovery after a network-node outage or initial power loss. Resources do not have to be relearned through a LOCATE broadcast search after a router failure. This reduces spikes of broadcasts that might otherwise occur when the APPN network is restarted.

Partial Directory Entries

The second technique that you can use to minimize the LOCATE flows in your APPN network is to define the resources in the local directory database by identifying the end node or network node where the particular resource is located.

The following is a sample configuration:

```
appn partner-lu-location CISCO.LU21
owning-cp CISCO.CP2
complete
```

The preceding example defines the location of an LU named CISCO.LU21 that is located with end node or network node CISCO.CP2. This command improves network performance by allowing directed Locate, instead of a broadcast. The disadvantage is that definitions must be created. To alleviate this definition problem, it may be possible to use partially specified names to define multiple resources.

The following is a sample configuration:

```
Sample configuration:
appn partner-lu-location CISCO.LU
owning-cp CISCO.CP2
wildcard
complete
```

The preceding example defines the location of all the LUs prefixed with the characters LU. Obviously, a naming convention is essential to the success of this type of node definition.

Central Directory Server (CDS)/Client

The third technique that you can use to minimize the LOCATE flows in your APPN network is to use the CDS/client function. The APPN architecture specifies a CDS that allows a designated network node to act as a focal point for locating network resources. In current APPN networks, every network node can potentially perform a broadcast search for a resource. This is because the directory services database is not replicated on every network node.

The CDS function allows a network node, with central directory client support, to send a directed LOCATE search to a CDS. If the CDS has no knowledge of the resource, it performs one broadcast search to find the resource. After the resource is found, the CDS caches the results in its directory. Subsequently, the CDS can provide the location of the resource to other network nodes without performing another broadcast search. The Cisco network-node implementation supports the central directory client function. VTAM is the only product that currently implements the CDS function.

Using the CDS means that there is a maximum of one broadcast search per resource in the network. This significantly reduces the amount of network traffic used for resource broadcast searching. You can define multiple CDSs in an APPN network. A network node learns the existence of a CDS via TDU exchange. If more than one CDS exists, the nearest one is used based on the number of hop counts. If a CDS fails, the route to the nearest alternative CDS is calculated automatically.

Central Resource Registration

The fourth technique that you can use to minimize the LOCATE flows in your APPN network is to use the central resource registration function. An end node registers its local resources at its network-node server. If every resource is registered, all network nodes can query the CDS, which eliminates the need for broadcast searches.

Backup Techniques in an APPN Network


This section provides an overview of the various backup techniques in APPN network. The backup and recovery scenarios are representative of common environments and requirements. The following three backup scenarios are discussed:

- A secondary WAN link as a backup to a primary WAN link
- Dual WAN links and dual routers providing full redundancy
- APPN DLUR backup support using a Cisco CIP router

Link Backup

The first backup technique that you can use in your APPN network is to use a secondary WAN link as a backup to your primary WAN link. By using the concept of auto-activation on demand, you can back up a primary WAN link with a secondary WAN link by using any supported protocols (for example, Point-to-Point [PPP], Switched Multimegabit Data Service [SMDS], and X.25), as shown in Figure 6-14.

Figure 6-14 Link backup.



In Figure 6-14, the Frame Relay link is the primary link and the ISDN dial link is the backup link. The requirement is that the ISDN link provides instantaneous backup for the primary link and it remains inactive until the primary link goes down. No manual intervention is needed. To support this, NNA needs to define two parallel transmission groups to NNB.

The primary link is defined using the following configuration command:

```
appn link-station PRIMARY
port FRAME_RELAY
fr-dest-address 35
retry-limit infinite
complete
```

The secondary link is defined as supporting auto-activation using the following configuration command:

```
appn link-station SECONDARY
port PPP
no connect-at-startup
adjacent-cp-name NETA.NNB
activate-on-demand
complete
```

By specifying **no connect-at-startup**, the secondary link is not activated upon APPN node startup. To indicate auto-activation support, specify **adjacent-cp-name** and **activate-on-demand**.

When the primary link fails, APPN detects the link failure and CP-CP sessions failure, which is disruptive to any existing LU-LU sessions. Because there are multiple links from NNA to NNB, NNA attempts to re-establish the CP-CP sessions over the secondary link. The CP-CP sessions request will activate the secondary dial link automatically.

To ensure that the Frame Relay link is used as primary and the dial PPP link is used as the backup, define the transmission group characteristics to reflect that. For example, use the **cost-per-connect-time** parameter to define the relative cost of using the dial PPP/ISDN link.


```
cost-per-connect-time 5
```

This will make the primary Frame Relay link a lower cost route. Therefore, it is a more desirable route than the secondary dial link because the default cost-per-connect-time is zero. When the primary link becomes active, there is no mechanism in place to automatically switch the sessions back to the primary link. Manual intervention is required.

Full Redundancy

The second backup technique that you can use in your APPN network is dual WAN links and dual routers for full redundancy. In some cases, for example, complete fault tolerance is required for mission-critical applications across the network. You can have dual routers and dual links installed to provide protection against any kind of communications failure.

Figure 6-15 shows how you can use duplicate virtual MAC addresses via RSRB to provide full redundancy and load sharing.

Figure 6-15 Full redundancy.

The router configuration for NNC is as follows:

```
source-bridge ring-group 200
!
interface TokenRing0
ring-speed 16
source 100 1 200
!
appn control-point NETA.NNC
complete
!
appn port RSRB rsrb
rsrb-virtual-station 4000.1000.2000 50 2 200
complete
```

The router configuration for NND is as follows:

```
source-bridge ring-group 300
!
interface TokenRing0
ring-speed 16
source 100 5 300
!
appn control-point NETA.NND
complete
!
appn port RSRB rsrb
rsrb-virtual-station 4000.1000.2000 60 3 300
complete
```

Both NNC and NND define an RSRB port with the same virtual MAC address. Every workstation will define the RSRB virtual MAC address as its destination MAC address of its network-node server. Essentially, a workstation can use either NNC or NND as its network-node server, depending on which node answers the test explorer frame first. The route to NNC consists of the following routing information:

```
Ring 100 -> Bridge 1 -> Ring 200 -> Bridge 2 -> Ring 50
```

Route to NND will consist of the following routing information:

```
Ring 100 -> Bridge 5 -> Ring 300 -> Bridge 3 -> Ring 60
```

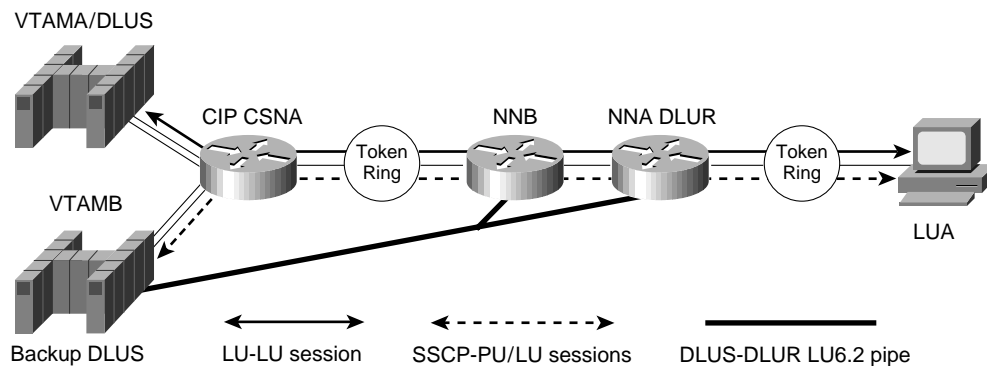
When NND fails, sessions on NND can be re-established over NNC instantaneously. This is analogous to the duplicate Token Ring interface coupler (TIC) support on the FEP except that no hardware is required. In Cisco's RSRB implementation, as shown in Figure 6-15, Segment 20 and Bridge 1, and Segment 30 and Bridge 2 are virtual. Duplicate MAC address can be supported without the hardware in place.

SSCP Takeover

The third backup technique is to use APPN DLUR with a Cisco CIP router to support transfer of resource ownership from one System Services Control Point (SSCP) (VTAM) to another when a failure occurs. This includes maintaining existing sessions over the failure. DLUS/DLUR can provide the capability to transfer SSCP ownership from the primary SSCP to the backup SSCP. It then examines how DLUR can provide the capability to obtain SSCP services from the backup SSCP without terminating LU-LU sessions that are in progress.

Figure 6-16 illustrates how the FEP can be replaced with a CIP router running CIP SNA (CSNA).

Figure 6-16 SSCP takeover with APPN and CIP.



In this example, VTAMA is the primary DLUS, VTAMB is the backup DLUS, and CIP router is configured as the DLUR. Assume that LUA requests to log on to an application that is residing on VTAMB. When VTAMA and the DLUS to DLUR connections fail, the DLUR node attempts to establish a session with VTAMB, which is configured as backup DLUS. When the control sessions to the DLUS are active, the DLUR node notifies VTAMB about all the active downstream physical and logical units. VTAMB sends active physical unit (ACTPU) and active logical unit (ACTLU) commands to these downstream devices. This transfers the resource ownership from VTAMA to VTAMB.

After the SSCP-PU and SSCP-LU sessions are re-established with VTAMB, new LU-LU sessions are possible. In addition, the DLUR node notifies VTAMB about all the dependent logical units that have active sessions.

The LU-LU path between VTAMB and LUA would be VTAMB -> NNB -> NNA -> LUA. When VTAMA fails, LU-LU sessions are not disrupted because VTAMA is not part of the LU-LU session path. In fact, LUA has no knowledge that the owning SSCP (VTAMA) failed and a new SSCP became the new owner. This process is transparent to LUA.

APPN in a Multiprotocol Environment

The trend in internetworking is to provide network designers with greater flexibility in building multiprotocol networks. Cisco provides the following two mechanisms to transport SNA traffic over an internetwork:

- Encapsulation
- Natively via APPN

The key to building multiprotocol internetworks is to implement some kind of traffic priority or bandwidth reservation to ensure acceptable response time for mission-critical traffic while maintaining some internetworking resource for less delay-sensitive traffic.

Bandwidth Management and Queuing

The following are some Cisco bandwidth management and queuing features that can enhance the overall performance of your network:

- Priority queuing
- Custom queuing
- Weighted fair queuing
- APPN buffer and memory management

For many years, the mainframe has been the dominant environment for processing business-critical applications. Increasingly powerful intelligent workstations, the creation of client-server computing environments, and higher bandwidth applications are changing network topologies. With the proliferation of LAN-based client-server applications, many corporate networks are migrating from purely hierarchical SNA-based networks to all-purpose multiprotocol internetworks that can accommodate the rapidly changing network requirements. This is not an easy transition. Network designers must understand how well the different protocols use shared network resources without causing excessive contentions among them.

Cisco has for many years provided technologies that encapsulate SNA traffic and allow consolidation of SNA with multiprotocol networks. APPN on the Cisco router provides an additional option in multiprotocol internetworks where SNA traffic can now flow natively and concurrently with other protocols. Regardless of the technology used in a multiprotocol environment, network performance is the key consideration.

Some of the major factors affecting network performance in a multiprotocol environment are as follows:

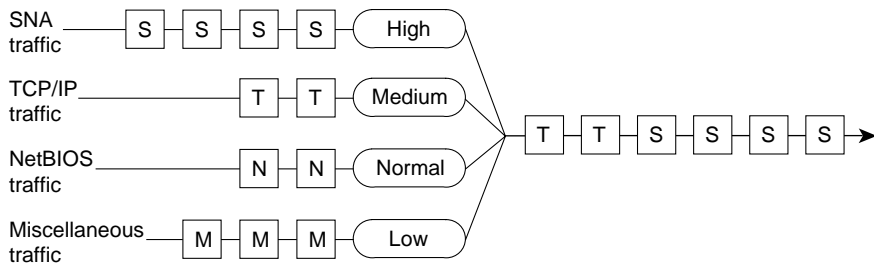
- *Media access speed*—The time it takes for a frame to be sent over a link. The capacity requirement of the network must be understood. Insufficient network capacity is the primary contributor to poor performance. Whether you have a single protocol network or a multiprotocol network, sufficient bandwidth is required.
- *Congestion control*—The router must have sufficient buffering capacity to handle instantaneous bursts of data. In order to support a multiprotocol environment, buffer management plays an important role to ensure that one protocol does not monopolize the buffer memory.
- *Latency in the intermediate routers*—This includes packet processing time while traversing a router and queuing delay. The former constitutes a minor part of the total delay. The latter is the major factor because client-server traffic is bursty.

Typically, subarea SNA traffic is highly predictable and has low bandwidth requirements. Compared to SNA traffic, client-server traffic tends to be bursty in nature and has high bandwidth requirements. Unless there is a mechanism in place to protect mission-critical SNA traffic, network performance could be impacted.

Cisco provides many internetworking solutions to enterprise networks by allowing the two types of traffic with different characteristics to coexist and share bandwidth; at the same time providing protection for mission-critical SNA data against less delay-sensitive client-server data. This is achieved through the use of several priority queuing and/or bandwidth reservation mechanisms.

For example, interface priority output queuing provides a way to prioritize packets transmitted on a per interface basis. The four possible queues associated with priority queuing—high, medium, normal and low—are shown in Figure 6-17. Priorities can be established based upon the protocol type, particular interface, SDLC address, and so forth.

Figure 6-17 Four queues of priority queuing.

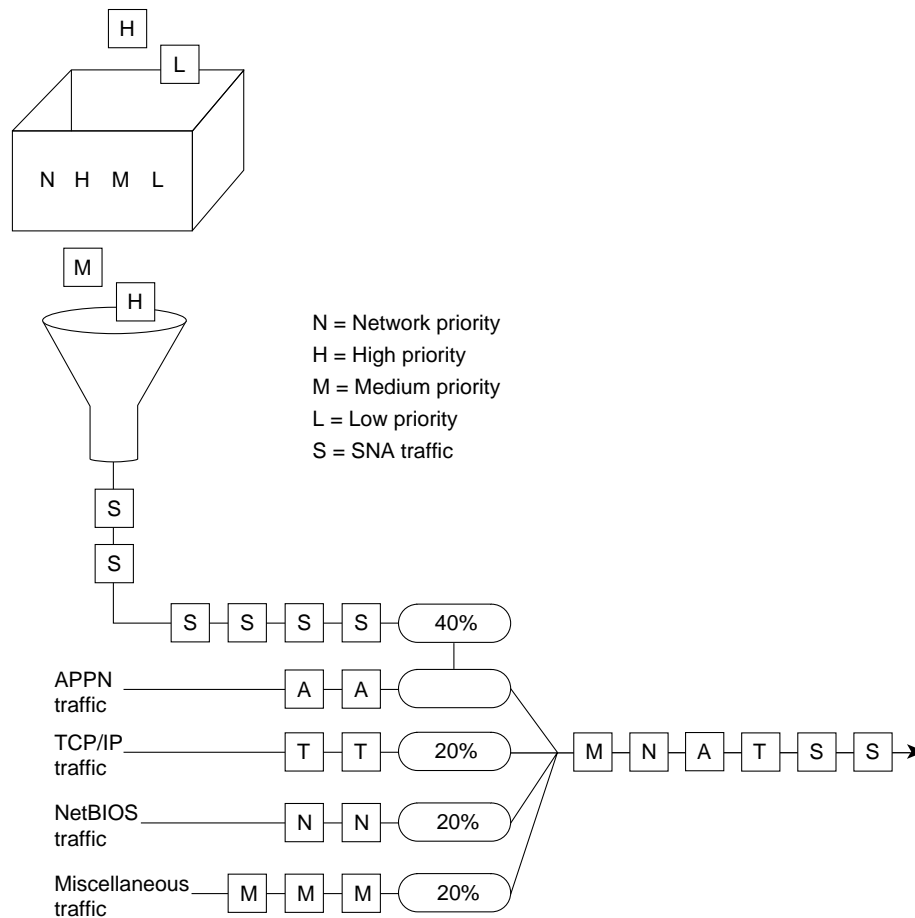


In Figure 6-18, SNA, TCP/IP, NetBIOS and other miscellaneous traffic are sharing the media. The SNA traffic is prioritized ahead of all other traffic, followed by TCP/IP, and then NetBIOS, and finally other miscellaneous traffic. There is no aging algorithm associated with this type of queuing. Packets that are queued to the high priority queue are always serviced prior to the medium queue, the medium queue is always serviced before the normal queue, and so forth.

Priority queuing, however, introduces a fairness problem in that packets classified to lower priority queues might not get serviced in a timely manner, or at all. Custom queuing is designed to address this problem. Custom queuing allows more granularity than priority queuing. In fact, this feature is commonly used in the internetworking environment in which multiple higher-layer protocols are supported. Custom queuing reserves bandwidth for a specific protocol, thus allowing mission-critical traffic to receive a guaranteed minimum amount of bandwidth at any time.

The intent is to reserve bandwidth for a particular type of traffic. For example, in Figure 6-18, SNA has 40 percent of the bandwidth reserved using custom queuing, TCP/IP 20 percent, NetBIOS 20 percent, and the remaining protocols 20 percent. The APPN protocol itself has the concept of COS that determines the transmission priority for every message. APPN prioritizes the traffic before sending it to the DLC transmission queue.

Figure 6-18 Example of custom queuing.



Custom queuing prioritizes multiprotocol traffic. A maximum of 16 queues can be built with custom queuing. Each queue is serviced sequentially until the number of bytes sent exceeds the configurable byte count or the queue is empty. One important function of custom queuing is that if SNA traffic uses only 20 percent of the link, the remaining 20 percent allocated to SNA can be shared by the other traffic.

Custom queuing is designed for environments that want to ensure a minimum level of service for all protocols. In today's multiprotocol internetwork environment, this important feature allows protocols of different characteristics to share the media. For an overview of how to use the other types of queuing to allow multiple protocols to coexist within a router, review Chapter 2, "Internetworking Design Basics."

Other Considerations with a Multiprotocol Environment

The memory requirement to support APPN is considerably higher than other protocols because of its large COS tables, network topology databases, and directory databases. To ensure that APPN will coexist with other network protocols when operating in a multiprotocol environment, users can define the maximum amount of memory available to APPN. The following is the sample configuration command.

```
apnn control-point CISCONET.EARTH
maximum-memory 16
complete
```

The preceding command specifies that APPN will not use more than 16 megabytes (MB) of memory. The memory is then managed locally by APPN. You can also specify the amount of memory reserved for APPN by using the following command:

```
appn control-point CISCONET.EARTH
  minimum-memory 32
complete
```

Note Memory that is dedicated to APPN is not available for other processing. Use this command with caution.

Although memory determines factors such as the number of sessions that APPN can support, buffer memory is required to regulate traffic sent to and from the router. To ensure that APPN has adequate buffers to support the traffic flows, you can define the percentage of buffer memory that is reserved for use by APPN. This prevents APPN from monopolizing the buffer memory available in the router.

The following is the sample configuration command.

```
appn control-point CISCONET.EARTH
  buffer-percent 60
complete
```

APPN uses a statistical buffering algorithm to manage the buffer usage. When buffer memory is constrained, APPN uses various flow control mechanisms to protect itself from severe congestion or deadlock conditions as a result of buffer shortage.

Network Management

As networks grow in size and complexity, there are many ways to provide network management for an enterprise. Table 6-2 summarizes Cisco’s management products.

Table 6-2 Network Management Tools Available for APPN Networks

Application	Description
Show commands	<p>A common challenge in APPN networks is to understand the topology and status of the resources in the network. Show commands take advantage of the fact that all network nodes in a network (or subnetwork) have a fully replicated network topology database. Only a single network node is required to get a view of the APPN subnet, and it should not matter which network node is chosen. In order to obtain more detailed information, such as attached end nodes and LEN nodes, and local ports and link stations, additional network nodes should be checked.</p> <p>The Cisco router supports the RFC1593, APPN MIB, which is used by the IBM 6611 router, so it can be an agent for SNMP APPN applications. Most APPN nodes can show much of this information in tabular form. In the Cisco router, the show appn topo command displays the topology database in tabular form. The show appn? command lists all of the options available.</p>
CiscoWorks Blue Maps	<p>A CiscoWorks application that shows logical maps of APPN, RSRB, and DLSw+ networks. It runs on the HP/UX, SunOS, and AIX operating systems. The APPN map is a manager for APPN SNMP agents, and displays the APPN network. The application can handle only a single network topology agent. If there are multiple subnets, the application can be started multiple times.</p>

Native Service Point (NSP)	<p>In SNA, a session between an SSCP and a PU is referred to as an SSCP-PU session. SSCPs use SSCP-PU sessions to send requests and receive status information from individual nodes. This information is then used to control the network configuration.</p> <p>NSP in the router can be used to send alerts and respond to requests from NetView on the mainframe computer. A service point allows NetView to establish a session to the router with the help of Cisco's applications that run on NetView. These applications cause commands to be sent to the router, and the router returns the reply. Currently, this is supported only over the SSCP-PU session, but DLUR can be used to accomplish this over an APPN network.</p>
Alerts and Traps	<p>NetView is the primary destination of alerts. It supports receiving alerts from both APPN and on the SSCP-PU session used by NSP. The Cisco router can send alerts on each session. At this time, two sessions are required: one for APPN-unique alerts and one for all other alerts. The new APPN MIB allows for APPN alerts to be sent as traps as well, with the Alert ID and affected resource included in the trap.</p> <p>To send alerts to NetView, the following command must be entered at NetView:</p> <pre style="margin-left: 40px;">FOCALPT CHANGE, FPCAT=ALERT, TARGET=NETA.ROUTER</pre>

Configuration Examples

This section provides the following APPN network configuration examples:

- Simple APPN network
- APPN network with end stations
- APPN over DLSw+

It also provides the following examples of using APPN when designing your network:

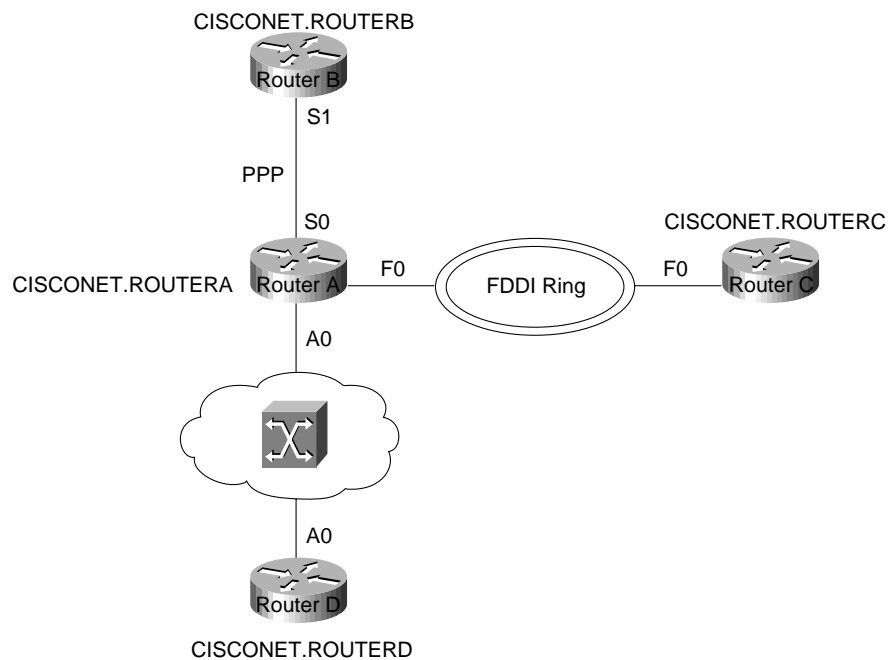
- Subarea to APPN migration
- APPN/CIP in a Sysplex environment
- APPN with FRAS BNN

As the following examples show, the minimal configuration for an APPN node includes an APPN control-point statement for the node and a port statement for each interface.

Simple APPN Network Configuration

Figure 6-19 shows an example of a simple APPN network that consists of four network nodes: Routers A, B, C, and D. Router A is responsible for initiating the connections to Routers B, C, and D. Consequently, it needs to define APPN logical links specifying the FDDI address of Router C, the ATM address of Router D, and so forth. For Routers B, C, and D, they can dynamically create the link-station definitions when Router A connects.

Figure 6-19 Example of a simple APPN network configuration.



Sample Configurations

This section provides sample configurations for each of these four network nodes (Routers A, B, C, and D) shown in Figure 6-19.

Router A Configuration

The following is a sample configuration for Router A shown in Figure 6-19. Note that all link stations are defined in Router A and dynamically discovered by the other routers. A link station connects two resources and must be defined with the destination address in one of the resources:

```

!
hostname routera
!
interface Serial0
 ip address 10.11.1.1 255.255.255.0
 encapsulation ppp
 no keepalive
 no fair-queue
 clockrate 4000000
!
interface Fddi0
 no ip address
 no keepalive
!
interface ATM0
 no ip address
 atm clock INTERNAL
 atm pvc 1 1 32 aal5nlpid
!
appn control-point CISCONET.ROUTERA
 complete
!
appn port PPP Serial0
 complete
!
appn port FDDI Fddi0
 desired-max-send-btu-size 3849
 max-rcv-btu-size 3849
 complete
!
appn port ATM ATM0
 complete
!
appn link-station LINKTOB
 port PPP
 complete
!
appn link-station LINKTOC
 port FDDI
 lan-dest-address 0000.6f85.a8a5
 no connect-at-startup
 retry-limit infinite 5
 complete
!
appn link-station LINKTOD
 port ATM
 atm-dest-address 1
 no connect-at-startup
 retry-limit infinite 5
 complete
!

```

Router B Configuration

The following is a sample configuration for Router B shown in Figure 6-19:

```
!  
hostname routerb  
!  
interface Serial1  
  ip address 10.11.1.2 255.255.255.0  
  encapsulation ppp  
  no keepalive  
  no fair-queue  
!  
appn control-point CISCONET.ROUTERB  
  complete  
!  
appn port PPP Serial1  
  complete  
!  
appn routing  
!  
end
```

Router C Configuration

The following is a sample configuration for Router C shown in Figure 6-19:

```
!  
hostname routerc  
!  
interface Fddi0  
  no ip address  
  no keepalive  
!  
appn control-point CISCONET.ROUTERC  
  complete  
!  
appn port FDDI Fddi0  
  desired-max-send-btu-size 3849  
  max-rcv-btu-size 3849  
  complete  
!  
appn routing  
!  
end
```

Router D Configuration


The following is a sample configuration for Router D shown in Figure 6-19:

```
!  
hostname routerd  
!  
interface ATM0  
  ip address 100.39.15.3 255.255.255.0  
  atm pvc 1 1 32 aal5nlpid  
!  
appn control-point CISCONET.ROUTERD  
  complete  
!  
appn port ATM ATM0  
  complete  
!  
appn routing  
!  
end
```

APPN Network Configuration with End Stations

Figure 6-20 shows an example of an APPN network with end stations. At the remote location, Router B initiates the APPN connection to Router A at the data center.

Figure 6-20 Example of an APPN network with end stations.



Sample Configurations

This section provides sample configurations for Routers A, B, and C shown in Figure 6-20.

Sample Configuration for Router A

The following is a sample configuration for Router A in Figure 6-20, which is responsible for initiating the APPN connection to the VTAM host:

```
hostname routera
!
interface TokenRing0
  no ip address
  mac-address 4000.1000.1000
  ring-speed 16
!
interface Serial0
  mtu 4096
  encapsulation frame-relay IETF
  keepalive 12
  frame-relay lmi-type ansi
  frame-relay map llc2 35
!
appn control-point CISCONET.ROUTERA
  complete
!
appn port FR0 Serial0
  complete
!
appn port TR0 TokenRing0
  complete
!
appn link-station TOVTAM
  port TR0
  lan-dest-address 4000.3745.0000
  complete
!
end
```

Sample Configuration for Router B

The following is a sample configuration for Router B shown in Figure 6-20. At the remote location, Router B initiates the APPN connection to Router A at the data center and EN AS/400. Because a link station is not defined in Router B for CISCONET.ENCM2B, a link station must be defined in ENCM2B for Router B:

```
!hostname routerb
!
interface TokenRing0
  mac-address 4000.1000.2000
  no ip address
  ring-speed 16
!
interface Serial0
  mtu 4096
  encapsulation frame-relay IETF
  keepalive 12
  frame-relay lmi-type ansi
  frame-relay map llc2 35
!
interface Serial2/7
  no ip address
  encapsulation sdlc
  no keepalive
  clockrate 19200
  sdlc role prim-xid-poll
  sdlc address 01
!
appn control-point CISCONET.ROUTERB
  complete
!
appn port FR0 Serial0
  complete
!
appn port SDLC Serial1
  sdlc-sec-addr 1
  complete
!
appn port TR0 TokenRing0
  complete
!
appn link-station AS400
  port SDLC
  role primary
  sdlc-dest-address 1
  complete
!
appn link-station ROUTERA
  port FR0
  fr-dest-address 35
  complete
!
end
```

Sample Configuration for Router C

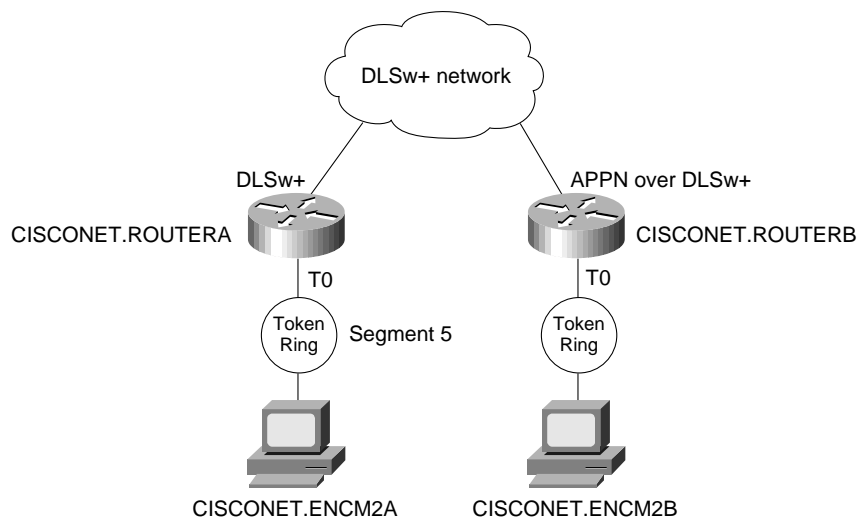
The following is a sample configuration for Router C shown in Figure 6-20. Router C initiates an APPN connection to Router A. Because there is not a link station for CISCONET.ENCM2C, one must be defined in the configuration for ENCM2C:

```
hostname routerc
!
interface TokenRing0
 mac-address 4000.1000.3000
 no ip address
 ring-speed 16
!
interface Serial0
 mtu 4096
 encapsulation frame-relay IETF
 keepalive 12
 frame-relay lmi-type ansi
 frame-relay map llc2 36
!
appn control-point CISCONET.ROUTERC
 complete
!
appn port FR0 Serial0
 complete
!
appn port TR0 TokenRing0
 complete
!
appn link-station TORTRA
 port FR0
 fr-dest-address 36
 complete
!
end
```

APPN over DLSw+ Configuration Example

Figure 6-21 shows an example of APPN with DLSw+. ROUTER A is a DLSw+ router with no APPN functions and ROUTERB is running DLSw+ and APPN.

Figure 6-21 Example of APPN with DLSw+.



Sample Configurations of DLSw+ Router A

The following section provides sample configurations for ROUTERA and ROUTERB and the two workstations shown in Figure 6-21.

Sample Configuration of DLSw+ ROUTERA

The following is a sample configuration for the DLSw+ ROUTERA shown in Figure 6-21:

```
hostname routera
!
source-bridge ring-group 100
dlsw local-peer peer-id 10.4.21.3
dlsw remote-peer 0 tcp 10.4.21.1
!
interface Serial0
  mtu 4096
  ip address 10.4.21.3 255.255.255.0
  encapsulation frame-relay IETF
  keepalive 12
  no fair-queue
  frame-relay lmi-type ansi
  frame-relay map llc2 56
!
interface TokenRing0
  ip address 10.4.22.2 255.255.255.0
  ring-speed 16
  multiring all
  source-bridge 5 1 100
!
```

Sample Configuration for Workstation Attached to ROUTERA

The following is a sample CS/2 configuration for the OS/2 workstation named CISCONET.ENCM2A shown in Figure 6-21. This workstation is attached to the DLSw+ router named ROUTERA. The workstation is configured as an end node and it uses ROUTERB as the

network-node server. The destination MAC address configured on this workstation is the virtual MAC address configured in ROUTERB on the **appn port** statement. A sample of the DLSw+ ROUTERB configuration is provided in the next section.

```

DEFINE_LOCAL_CP  FQ_CP_NAME ( CISCONET. ENCM2A )
                  CP_ALIAS ( ENCM2C )
                  NAU_ADDRESS ( INDEPENDENT_LU )
                  NODE_TYPE ( EN )
                  NODE_ID ( X' 05D00000 ' )
                  NW_FP_SUPPORT ( NONE )
                  HOST_FP_SUPPORT ( YES )
                  MAX_COMP_LEVEL ( NONE )
                  MAX_COMP_TOKENS ( 0 ) ;
DEFINE_LOGICAL_LINK  LINK_NAME ( TORTRB )
                    ADJACENT_NODE_TYPE ( LEARN )
                    PREFERRED_NN_SERVER ( YES )
                    DLC_NAME ( IBMTRNET )
                    ADAPTER_NUMBER ( 0 )
                    DESTINATION_ADDRESS ( X' 400010001112 ' )
                    ETHERNET_FORMAT ( NO )
                    CP_CP_SESSION_SUPPORT ( YES )
                    SOLICIT_SSCP_SESSION ( YES )
                    NODE_ID ( X' 05D00000 ' )
                    ACTIVATE_AT_STARTUP ( YES )
                    USE_PUNAME_AS_CPNAME ( NO )
                    LIMITED_RESOURCE ( NO )
                    LINK_STATION_ROLE ( USE_ADAPTER_DEFINITION )
                    MAX_ACTIVATION_ATTEMPTS ( USE_ADAPTER_DEFINITION )
                    EFFECTIVE_CAPACITY ( USE_ADAPTER_DEFINITION )
                    COST_PER_CONNECT_TIME ( USE_ADAPTER_DEFINITION )
                    COST_PER_BYTE ( USE_ADAPTER_DEFINITION )
                    SECURITY ( USE_ADAPTER_DEFINITION )
                    PROPAGATION_DELAY ( USE_ADAPTER_DEFINITION )
                    USER_DEFINED_1 ( USE_ADAPTER_DEFINITION )
                    USER_DEFINED_2 ( USE_ADAPTER_DEFINITION )
                    USER_DEFINED_3 ( USE_ADAPTER_DEFINITION ) ;
DEFINE_DEFAULTS  IMPLICIT_INBOUND_PLU_SUPPORT ( YES )
                 DEFAULT_MODE_NAME ( BLANK )
                 MAX_MC_LL_SEND_SIZE ( 32767 )
                 DIRECTORY_FOR_INBOUND_ATTACHES ( * )
                 DEFAULT_TP_OPERATION ( NONQUEUED_AM_STARTED )
                 DEFAULT_TP_PROGRAM_TYPE ( BACKGROUND )
                 DEFAULT_TP_CONV_SECURITY_RQD ( NO )
                 MAX_HELD_ALERTS ( 10 ) ;
START_ATTACH_MANAGER ;

```

Sample Configuration for DLSw+ ROUTERB

ROUTERB, shown in Figure 6-21, is an APPN router that uses the APPN over DLSw+ feature. The VDLC operand on the port statement indicates that APPN is carried over DLSw+. The following is a sample configuration for this router:

```
hostname routerb
!
source-bridge ring-group 100
dlsw local-peer peer-id 10.4.21.1
dlsw remote-peer 0 tcp 10.4.21.3
!
interface Serial2/0
mtu 4096
ip address 10.4.21.1 255.255.255.0
encapsulation frame-relay IETF
keepalive 12
no fair-queue
frame-relay map llc2 35
!
interface TokenRing0
no ip address
ring-speed 16
mac-address 4000.5000.6000
!
appn control-point CISCONET.ROUTERB
complete
!
appn port VDLC vdlc
vdlc 100 vmac 4000.1000.1112
complete
!
```

Sample Configuration for Workstation Attached to ROUTERB

The following is a sample CS/2 configuration for the OS/2 workstation named CISCONET.ENCM2B shown in Figure 6-21. This workstation is attached to the DLSw+ router named ROUTERB:

```

DEFINE_LOCAL_CP  FQ_CP_NAME(CISCONET.ENCM2B)
                  CP_ALIAS(ENCM2C)
                  NAU_ADDRESS(INDEPENDENT_LU)
                  NODE_TYPE(EN)
                  NODE_ID(X'05D00000')
                  NW_FP_SUPPORT(NONE)
                  HOST_FP_SUPPORT(YES)
                  MAX_COMP_LEVEL(NONE)
                  MAX_COMP_TOKENS(0);

DEFINE_LOGICAL_LINK LINK_NAME(TORTRB)
                    ADJACENT_NODE_TYPE(LEARN)
                    PREFERRED_NN_SERVER(YES)
                    DLC_NAME(IBMTRNET)
                    ADAPTER_NUMBER(0)
                    DESTINATION_ADDRESS(X'400050006000')
                    ETHERNET_FORMAT(NO)
                    CP_CP_SESSION_SUPPORT(YES)
                    SOLICIT_SSCP_SESSION(YES)
                    NODE_ID(X'05D00000')
                    ACTIVATE_AT_STARTUP(YES)
                    USE_PUNAME_AS_CPNAME(NO)
                    LIMITED_RESOURCE(NO)
                    LINK_STATION_ROLE(USE_ADAPTER_DEFINITION)
                    MAX_ACTIVATION_ATTEMPTS(USE_ADAPTER_DEFINITION)
                    EFFECTIVE_CAPACITY(USE_ADAPTER_DEFINITION)
                    COST_PER_CONNECT_TIME(USE_ADAPTER_DEFINITION)
                    COST_PER_BYTE(USE_ADAPTER_DEFINITION)
                    SECURITY(USE_ADAPTER_DEFINITION)
                    PROPAGATION_DELAY(USE_ADAPTER_DEFINITION)
                    USER_DEFINED_1(USE_ADAPTER_DEFINITION)
                    USER_DEFINED_2(USE_ADAPTER_DEFINITION)
                    USER_DEFINED_3(USE_ADAPTER_DEFINITION);

DEFINE_DEFAULTS  IMPLICIT_INBOUND_PLU_SUPPORT(YES)
                  DEFAULT_MODE_NAME(BLANK)
                  MAX_MC_LL_SEND_SIZE(32767)
                  DIRECTORY_FOR_INBOUND_ATTACHES(*)
                  DEFAULT_TP_OPERATION(NONQUEUED_AM_STARTED)
                  DEFAULT_TP_PROGRAM_TYPE(BACKGROUND)
                  DEFAULT_TP_CONV_SECURITY_RQD(NO)
                  MAX_HELD_ALERTS(10);

START_ATTACH_MANAGER;

```


Note For more information on DLSw+, see Chapter 7, “Designing DLSw+ Internetworks.”

Example of Subarea to APPN Migration

This section provides an overview of the implementation and conversion of the SNA network from subarea FEP-based to APPN router-based. It explores the use of DLSw+ as a migration technology from traditional SNA to APPN, and covers the migration steps. The example involves a large insurance company in Europe. The company plans to replace the FEPs with Cisco routers, migrating from subarea to APPN routing.

Figure 6-22 shows the company's current SNA network. The network consists of two mainframe sites running four VTAM images with a Communications Management Complex (CMC) host in each data center, as shown in Figure 6-22. In every data center, four NCR Comten FEPs (IBM 3745-compatible) support traffic from multiple regional offices. There are also two NCR Comten FEPs that provide SNA Network Interconnect (SNI) support.

Figure 6-22 SNA FEP-based network.

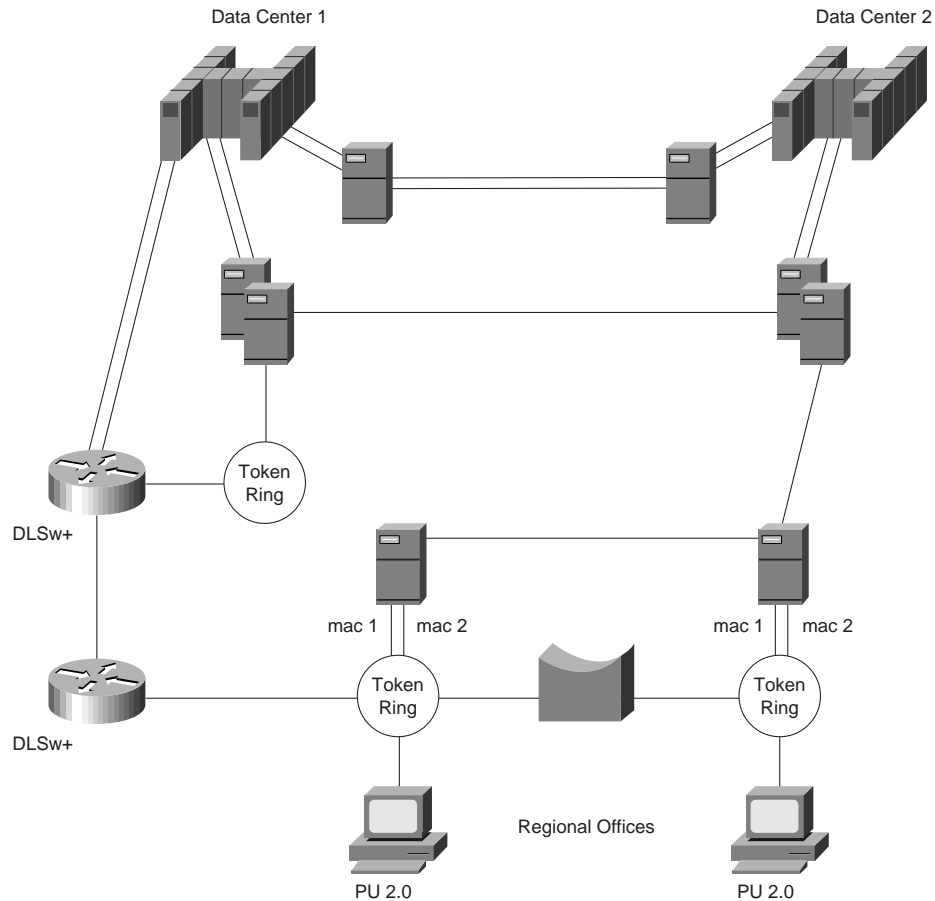


There are 22 regional offices across the country. Every regional office has two NCR Comten FEPs installed, one connecting to Data Center 1 and the other connecting to Data Center 2. The remote FEPs have dual Token Rings that are connected via a bridge; duplicate TIC address support is implemented for backup and redundancy. This means that a PU2.0 station can connect to the host through any one of the two FEPs. If one FEP fails, PU2.0 stations can access the host via the other FEP.

In addition to the Token-Ring-attached devices (approximately 15 per regional office), the two FEPs also run NCP Packet-Switching Interface (NPSI), supporting over 200 remotely attached devices via the public X.25 network. The total number of LUs supported per regional office is approximately 1800, with 1500 active LU-LU sessions at any one time. The estimated traffic rate is 15 transactions per second.

The first migration step is to implement Cisco CIP routers at one of the data centers, replacing the channel-attached FEPs. A remote router is then installed in one of the regional offices. The two routers are connected using DLSw+, as shown in Figure 6-23.


Figure 6-23 Subarea to APPN migration—phase one.



As Figure 6-23 shows, the FEPs at the regional office continue to provide boundary functions to the Token Ring and X.25-attached devices. The two DLSw+ routers handle the traffic between the FEP at Data Center 1 and the FEP at the regional office. SNA COS is preserved in this environment.

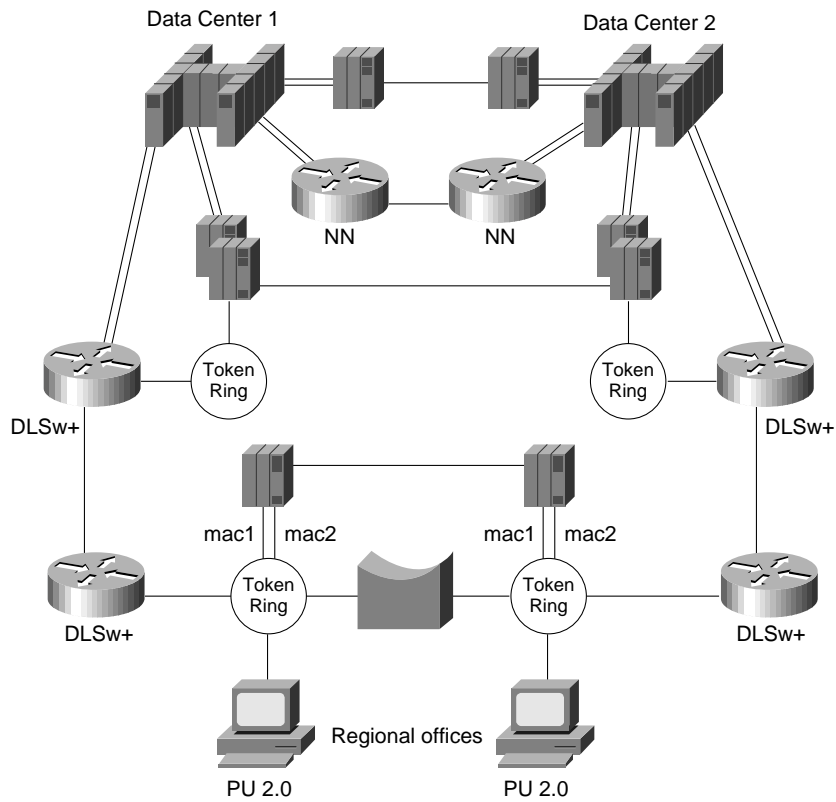
After stability of the routers is ensured, the network designer proceeds to the next phase. As Figure 6-24 shows, this phase involves installation of a second router in Data Center 2 and the regional office. At this point, FEP-to-FEP communications between regional offices and data centers are handled by the routers via DLSw+.

Figure 6-24 Subarea to APPN migration—phase two.



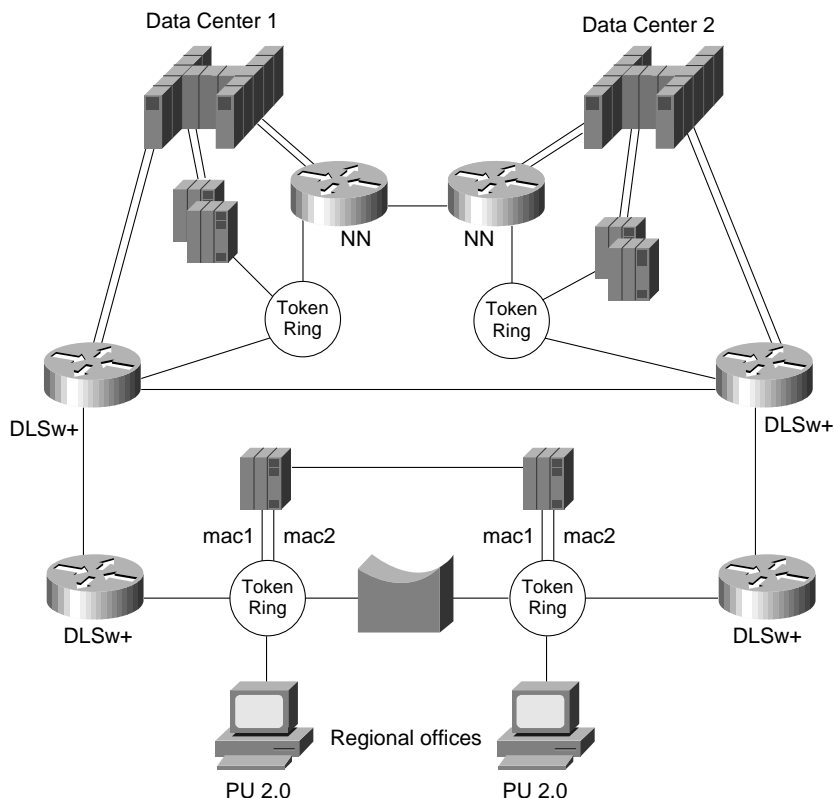
Continuing with the migration plan, the network designer's next step is to install an additional CIP router in each data center to support traffic between the two data centers. As shown in Figure 6-25, the links that are connecting the FEPs in Data Center 1 and Data Center 2 are moved one by one to the routers.

Figure 6-25 Subarea to APPN migration—phase three.



APPN will be enabled to support the traffic between Data Center 1 and Data Center 2. Eventually, the FEP-based network will become a router-based network. The NCR Comten processors will become obsolete. Two of the NCR Comten processors will be kept to provide SNI support to external organizations. Figure 6-26 illustrates the new router-based network.


Figure 6-26 Subarea to APPN migration—phase four.



The communication links that formerly connected the FEPs in the two data centers are now moved to the routers. The FEPs at the data centers can be eliminated. The FEPs at the regional offices are merely providing the boundary functions for dependent LU devices, thus allowing SNA COS to be maintained. The next phase is to migrate the SNA boundary functions support from the FEP to the remote router at the regional office by enabling APPN and DLUR. After this is complete, all the FEPs can be eliminated.

The next step is to migrate from DLSw+ to APPN between the data center routers and the regional office routers. This is done region by region until stability of the network is ensured. As shown in Figure 6-27, DLUR is enabled to support the dependent PU devices in the regional offices. X.25 attached dependent PU2.0 devices that are formerly connected to the FEPs using NPSI are supported via Qualified Logical Link Control (QLLC) in the router. QLLC is the standard for SNA encapsulation for X.25.

Figure 6-27 Subarea to APPN migration—phase five.



Example of APPN/CIP in a Sysplex Environment

This section examines APPN and the CIP routers in a Sysplex (system complex) environment. It provides an overview of the Sysplex environment and its relationship with APPN along with a description of how to use the following three approaches to support the Sysplex environment:

- Sysplex with APPN Using Subarea Routing—Option One
- Sysplex Using Subarea/APPN Routing—Option Two
- Sysplex Using APPN Routing—Option Three

It also describes how APPN provides fault tolerance and load sharing capabilities in the data center.

Sysplex Overview

Sysplex provides a means to centrally operate and manage a group of multiple virtual storage (MVS) systems by coupling hardware elements and software services. Many data processing centers have multiple MVS systems to support their business, and these systems often share data and applications. Sysplex is designed to provide a cost-effective solution to meet a company's expanding requirements by allowing MVS systems to be added and managed efficiently.

A Sysplex environment consists of multiple 9672 CMOS processors, and each CMOS processor presents a VTAM domain. The concept of multiprocessors introduces a problem. Today, users are accustomed to single images. For example, IMS (Information Management System) running on the

mainframe can serve the entire organization on a single host image. With the multiprocessor concept, you would not want to instruct User A to establish the session with IMS on System A and User B to establish the session with IMS on System B because IMS might run on either system.


To resolve this, a function called *generic resource* was created. The generic resource function enables multiple application programs, which provide the same function, to be known and accessed by a single generic name. This means that User A might sometimes get IMS on System A, and sometimes get IMS on System B. Because both systems have access to the same shared data in the Sysplex, this switching of systems is transparent to the users. VTAM is responsible for resolving the generic name and determining which application program is used to establish the session. This function enables VTAM to provide workload balancing by distributing incoming session initiations among a number of identical application programs that are running on different processors.

Generic resource runs only on VTAM with APPN support. In order to achieve session load balancing across the different processors, users must migrate VTAM from subarea SNA to APPN. The rest of this section examines three options for supporting the Sysplex environment.

Sysplex with APPN Using Subarea Routing—Option One

The first option to support the Sysplex environment is to convert the CMC host to a composite network node. Traditionally, the CMC host was the VTAM that owned all of the network's SNA resources. With this approach, the composite network node is used to describe the combination of VTAM and Network Control Program (NCP). This means that VTAM and NCP function together as a single network node. In Figure 6-28, the CMC host and the FEPs are configured as the composite network node.

Figure 6-28 CMC composite network node with subarea routing—option one.



The VTAM CMC host owns the FEPs. Each FEP is connected to the 9672 CMOS processors through a parallel channel. Each 9672 CMOS processor is configured as a migration data host and maintains both an APPN and subarea appearance.


Each migration data host establishes subarea connections to the FEPs using Virtual Route Transmission Group (VRTG), which allows APPN to be transported over traditional subarea routing. CP-CP sessions between the CMC host and the 9672 migration data hosts are established using VRTG. Generic resource function is performed in APPN, but all routing is subarea routing. This is the most conservative way to migrate to a Sysplex.

The disadvantage of this approach is that using subarea routing does not provide dynamic implementation of topology changes in APPN, which is available with APPN connection. If you need to add a CMOS processor, subarea PATH changes to every subarea node are required. Another drawback of this approach is that running APPN over subarea routing introduces complexity to your network.

Sysplex Using Subarea/APPN Routing—Option Two

The second option to support the Sysplex environment is to use subarea/APPN routing. This approach is similar to Option One, which was described in the preceding section. With this second approach, the CMC host and the FEPs are converted to a composite network node, as shown in Figure 6-29.

Figure 6-29 CMC composite network node with APPN routing—option two.




As shown in Figure 6-29, the two 9672 CMOS processors are converted to pure end nodes (EN A and EN B). APPN connections are established between the 9672s and the FEPs. Sessions come into the CMC in the usual way and the CMC does subarea/APPN interchange function. This means that sessions are converted from subarea routing to APPN routing on the links between the FEPs and the 9672s.

A disadvantage of this second approach is that it performs poorly because the FEPs must perform an extra conversion. This approach also requires more NCP cycles and memory. Although this is very easy to configure and it does not require any changes to the basic subarea routing, the cost of the NCP upgrades can be expensive.

Sysplex Using APPN Routing—Option Three

The third option to support the Sysplex environment is to use APPN routing. With this approach, you use DLUR as a front end to the CMC-owned logical units. Figure 6-30 illustrates this configuration.

Figure 6-30 Sysplex with DLUR using CIP—option three.



As shown in Figure 6-30, this is a pure APPN network with APPN routing only. Each CMOS end-node processor is attached to the DLUR routers through APPN. Note that the DLUR routers could be remote and not directly next to the mainframe computers (for example, there could be intervening routers).

This is the preferred approach for implementing the Sysplex environment for the company used in this sample scenario. The following section provides more details on this sample implementation.

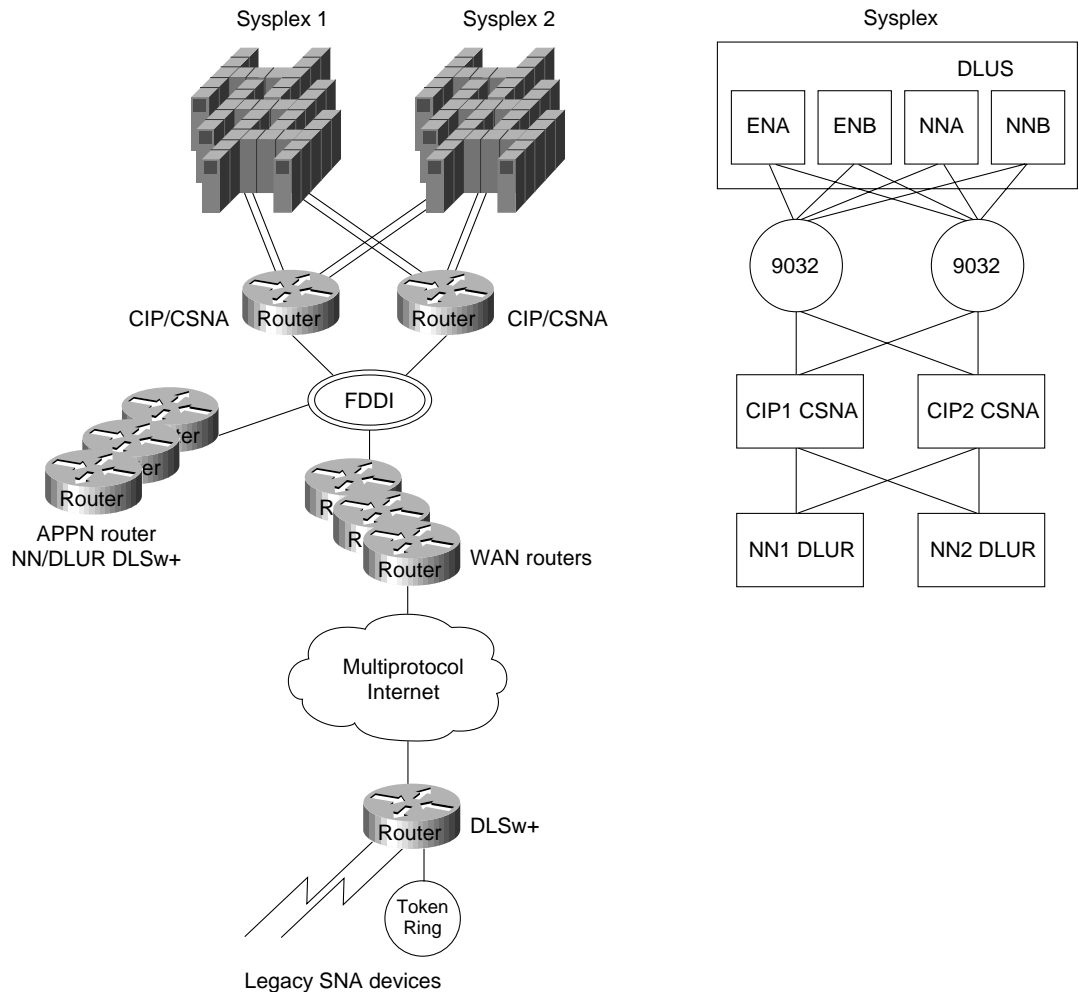
The Company's Network

The company used in this example has a very large IP backbone and a very large SNA network. Today, its multiprotocol and SNA network are separate. The company's goal is to consolidate the traffic across the multiprotocol Internet. The company has chosen IP as its strategic backbone protocol of choice. To transport the SNA traffic, DLSw+ is used.

In the data center, the company plans to support five different IBM Sysplex environments. Its objective is to have the highest degree of redundancy and fault tolerance. The administrators decided not to run APPN throughout their existing multiprotocol network but chose APPN in the data center to provide the required level of redundancy.

Figure 6-31 shows the configuration of the company's data center. The diagram on the top right in this figure is a logical view of one Sysplex environment and how it is connected to the multiprotocol network through the CIP/CSNA routers and the APPN routers. Each CIP/CSNA router has two parallel channel adapters to each Sysplex host (Sysplex 1 and Sysplex 2) through separate ESCON Directors. To meet the company's high availability requirement, this configuration has no single points of failure.

Figure 6-31 Example of APPN in the data center.



In each Sysplex environment, there are a minimum of two network nodes per Sysplex acting as a DLUS. VTAM NNA is designated as the primary DLUS node. NNB is designated the backup DLUS. The remaining hosts are data hosts configured as end nodes. These end node data hosts use NNA as the network-node server.

There are two CIP routers to support every Sysplex environment and at least two APPN routers running DLUR to provide boundary functions support for remote devices. The traffic is expected to load share across the two CIP routers. Consequently, APPN provides load balancing and redundancy in this environment.

Sample Configuration

From an APPN standpoint, NNA in Figure 6-31 can be configured as the primary DLUS. NNB can be configured as the backup DLUS. The following is a configuration example for NN1. NN2 would be configured similarly.

```

!
appn control-point CISCONET.NN1
  dlus CISCONET.NNA
  backup-dlus CISCONET.NNB
  dlur
complete
    
```

When the primary DLUS host goes out of service for any reason, the DLUR node is disconnected from its serving DLUS. The DLUR node retries the DLUS/DLUR pipe with NNA. If unsuccessful, it tries its backup DLUS.

To achieve load balancing, every DLUR router defines two parallel APPN transmission groups with equal weights to every VTAM host using the following configuration:

```
!
! Link to VTAM ENA via CIP router 1
!
appn link-station LINK1ENA
  port FDDI0
  lan-dest-address 4000.3000.1001
  complete
!
! Link to VTAM ENA via CIP router 2
!
appn link-station LINK2ENA
  port FDDI0
  lan-dest-address 4000.3000.2001
  complete
!
! Link to VTAM ENB via CIP router 1
!
appn link-station LINK1ENB
  port FDDI0
  lan-dest-address 4000.3000.1002
  complete
!
! Link to VTAM ENB via CIP router 2
!
appn link-station LINK2ENB
  port FDDI0
  lan-dest-address 4000.3000.2002
  complete
!
! Link to Primary DLUS NNA via CIP router 1
!
appn link-station LINK1NNA
  port FDDI0
  lan-dest-address 4000.3000.1003
  complete
!
! Link to Primary DLUS NNA via CIP router 2
!
appn link-station LINK2NNA
  port FDDI0
  lan-dest-address 4000.3000.2003
  complete
!
! Link to Backup DLUS NNB via CIP router 1
!
appn link-station LINK1NNB
  port FDDI0
  lan-dest-address 4000.3000.1004
  complete
!
! Link to Backup DLUS NNB via CIP router 2
!
appn link-station LINK2NNB
  port FDDI0
  lan-dest-address 4000.3000.2004
  complete
```

As shown in the preceding configuration, NN1 defines two APPN transmission groups to ENA, ENB, NNA and NNB. There are two channel attachments to each host and each attachment is connected to separate hardware (for example, a CIP card, CIP router, ESCON Director). Reasons to have duplicate hardware include provision for the loss of any physical component; if this happens the host is still accessible using the alternative path.

From an APPN perspective, there are two transmission groups that connect a DLUR router and every host. One transmission group traverses CIP Router 1 and the other traverses CIP Router 2. When one path fails, the APPN transmission group becomes inoperative. The second transmission group provides an alternative route for host connection through the other path.

All the subarea SSCP/PU and SSCP/LU sessions flow on one of the transmission groups between the DLUR router and the primary DLUS host. As for the LU-LU sessions, the two possible routes between the DLUR router and a VTAM host are available. The DLUR router and a VTAM host select one of these two routes at random for the LU-LU sessions. This randomization provides a certain amount of load distribution across the two CIP routers, although it might not necessarily be statistically load balanced.

There are multiple DLUR routers that support downstream SNA devices. The following is a sample configuration for DLUR router NN1:

```
source-bridge ring-group 100
dlsw local-peer peer-id 172.18.3.111 promiscuous
!
interface FDDI0
 ip address 172.18.3.111 255.255.255.0
!
appn control-point NETA.NN1
 complete
!
appn port VDLC1 vdlc
 vdlc 100 4000.1000.2000
 complete
```

The following is a sample configuration for DLUR router NN2:

```
source-bridge ring-group 200
dlsw local-peer peer-id 172.18.3.112 promiscuous
!
interface FDDI0
 ip address 172.18.3.112 255.255.255.0
!
appn control-point NETA.NN2
 complete
!
appn port VDLC2 vdlc
 vdlc 200 4000.1000.2000
 complete
```

A workstation gains access to the host through the DLUR router. A workstation defines 4000.1000.2000 as the destination MAC address in the emulation software. This virtual MAC address is defined to every DLUR router. When initiating a connection, a workstation sends an all-routes broadcast Test command frame to the MAC address to which it wants to connect. The remote DLSw+ router sends an explorer frame to its peers. Both NN1 and NN2 respond with **ICANREACH**. The DLSw+ router is configured to use the load balancing mode. This means that the DLSw+ router caches both NN1 and NN2 as peers that can reach the host. Host sessions are established through NN1 and NN2 in a round robin fashion. This allows the company to spread its SNA traffic over two or more DLUR routers. If NN1 becomes unavailable, sessions that traverse NN1 are disruptive but they can be re-established through NN2 with negligible impact.

This design increases overall availability by using duplicate virtual MAC address on the DLUR router. The dual paths provide the option for a secondary path to be available for use when the primary path is unavailable. Another advantage is that this design allows for easy scaling. For example, when the number of SNA devices increases, buffer memory might become a constraint on the DLUR routers. The company can add a DLUR router to support the increased session load. This topology change does not require any network administration from any remote routers or the data center routers.


Example of APPN with FRAS BNN

This section describes the design considerations when building a large enterprise APPN network. It lists the current technologies that allow the company in this example to build a large APPN network. Each option is discussed in detail. FRAS BNN is chosen as an interim scalability solution to reduce the number of network nodes in the network. This allows the network to scale to meet the company's expanding requirements.

In this example, a government agency has a network that consists of one data center and approximately 100 remote sites. Within the next few years, its network is expected to increase to 500 remote sites.

Figure 6-32 shows a simplified version of the agency's current APPN network.

Figure 6-32 Sample APPN network for a government agency.



The data center consists of 20 mainframe processors from IBM and a variety of other vendors. The IBM mainframes are MVS-based and are running VTAM. They are also configured as NN/DLUS and EN data hosts. No subarea protocol exists in this network. Other non-IBM mainframes are configured as either an EN or LEN node.

The user platform is OS/2 running Communications Server at all the remote sites with connectivity needs to the data center mainframe computers. Initially, there are no any-to-any communication requirements in this network. The applications supported are LU type 2 and LU6.2.

APPN in the Data Center

The host mainframes in Figure 6-32 are connected using the external communication adapter (XCA) connection over the 3172 Interconnect Controllers. The non-IBM data hosts (Companies A, B, and C) use the VTAM IBM mainframe as the network-node server. To keep the amount of TDU flows to a minimum, CP-CP sessions exist only between VTAM and the data center routers. There are no CP-CP sessions among the routers located at the data center.

To achieve the optimal route calculation without explicitly defining meshed connection definitions, every end node and network node at the data center is connected to the same connection network. This allows a session to be directly established between two data center resources without traversing the VTAM network node. As Figure 6-32 shows, when an LU-LU session between resources at EN3A and Company A's mainframe is set up, the optimal route is directly through the FDDI ring to NN1 and NN3.

To reduce the number of broadcast searches to a maximum of one per resource, VTAM is configured as the CDS in this network. The CDS function is very effective in this network because the resources in the network require access only to resources at the host mainframes in the data center. These host mainframes register their resources with VTAM, which is their network-node server. Consequently, VTAM always has location information for every resource at the data center. This means that VTAM never has to perform LOCATE broadcast searches.

APPN in the Remote Site

The network depicted in Figure 6-32 has approximately 30 to 40 CS/2 workstations in every remote site. Every user workstation is configured as an end node. Each end node supports eight independent LU6.2 sessions and four dependent LU sessions. A Cisco router at every location forwards the traffic to the data center. The router's network node function provides the intermediate routing node function for the independent LUs. The DLUR function provides the dependent LU routing function for the dependent LUs.

Future Configuration


This network will eventually consist of 500 remote network node routers, 100 data center routers, and eight mainframe computers. Typically, a 600-node APPN network will have scalability issues. The rest of this section examines the following two options that you can use to address scalability issues in an APPN network:

- Implementing border node on VTAM to partition the network into smaller subnets
- Using FRAS BNN to reduce the number of network nodes in the network

Using Border Node on VTAM to Partition the Network into Smaller Subnets

By implementing the concept of border node on VTAM, a peripheral subnetwork boundary is introduced between NN1 and VTAM, and between NN2 and VTAM, as shown in Figure 6-33.

Figure 6-33 APPN network with VTAM extended border node.



There would be no topology information exchange between VTAM and the data center NN routers. The eight mainframe computers would be in the same subnet. Every data center router would support multiple access routers and they would form their own subnet. Each subnet is limited to a maximum of 100 network nodes. This configuration would prevent topology information from being sent from one subnet to another, thus allowing the network to scale to over 600 network nodes.


Although this approach addresses the TDU flow issue, there is a considerable loss of functions, however, by configuring VTAM as a border node in this environment. First, two APPN subnetworks cannot be connected through a connection network. LU-LU sessions between resources at Company A's host and remote resources would be set up through an indirect route through the VTAM border node. This is clearly not an optimal route. Second, the central directory server function is lost because the VTAM border node portrays an end node image to NN1. This prevents NN1 from discovering the central directory server in the network.

The next section examines an alternative approach of using FRAS BNN to reduce the number of network nodes in the network.

Using FRAS BNN to Reduce the Number of Network Nodes

Figure 6-34 shows how FRAS BNN can be used to reduce the number of network nodes in the company's network. All the server applications are on the mainframe computers and devices only require host access. APPN routing is not essential for this company.

Figure 6-34 APPN network with FRAS BNN.



Implementing FRAS BNN rather than a full APPN network node on the access routers directly reduces the number of network nodes. This allows the network to scale without the concern of TDU flows. This is proven to be a viable solution for this company for the time being because LAN-to-LAN connectivity is not an immediate requirement. The remote routers can be migrated to support APPN border node when it becomes available.

In this environment, CP-CP sessions are supported over the Frame Relay network. The central directory server and the concept of Connection Network are fully supported. LU-LU sessions can be set up using the direct route without traversing VTAM, as shown in Figure 6-34. The only function that is lost with FRAS BNN is COS for traffic traveling from the remote FRAS BNN router to the data center.

Summary

Recall that this chapter discussed developing the network design and planning a successful migration to APPN. It covered the following topics:

- Evolution of SNA
- When to Use APPN as Part of a Network Design
- When to Use APPN Versus Alternative Methods of SNA Transport
- Overview of APPN

- Scalability Issues
- Backup Techniques in an APPN Network
- APPN in a Multiprotocol Environment
- Network Management
- Configuration Examples