

# Implementation of SNASw

## Getting Started

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

## When Is APPN Required?

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

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

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

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

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

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

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

## SNASw General Considerations

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

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

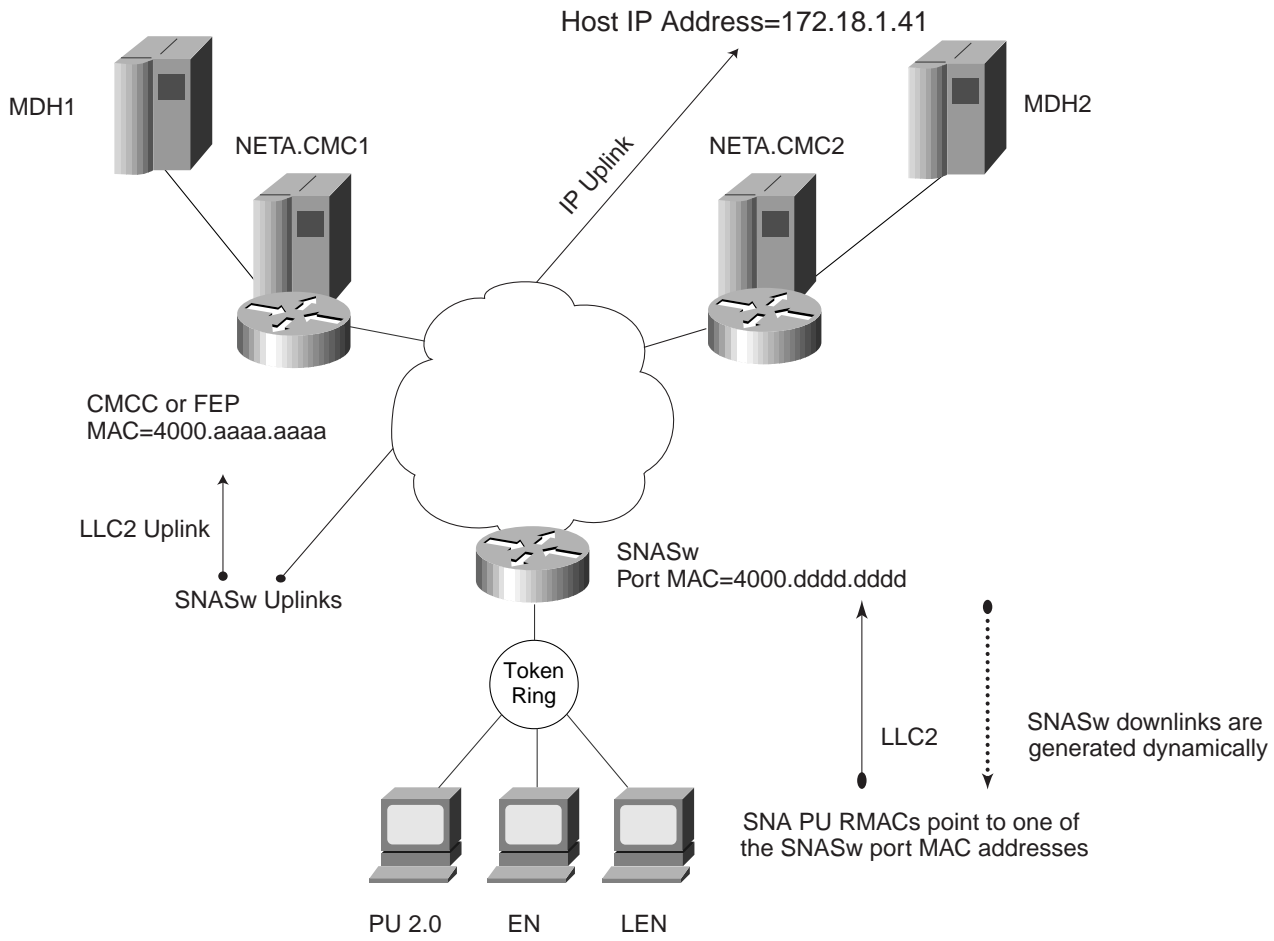
## Required SNASw Configuration

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

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

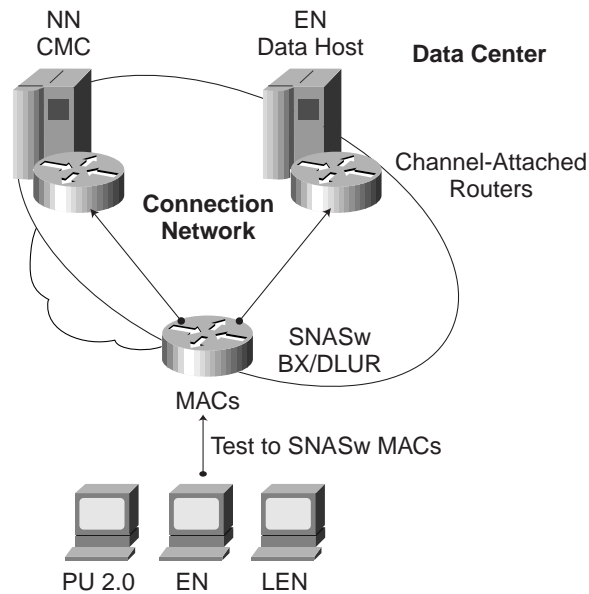
Note: SNASw supersedes all functionality previously available in the APPN NN feature in the Cisco IOS Software. SNASw configuration does not accept the previous APPN configuration commands.

Figure 3-1 Required Configuration



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

Figure 3-2 SNASw and Connection Network



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

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

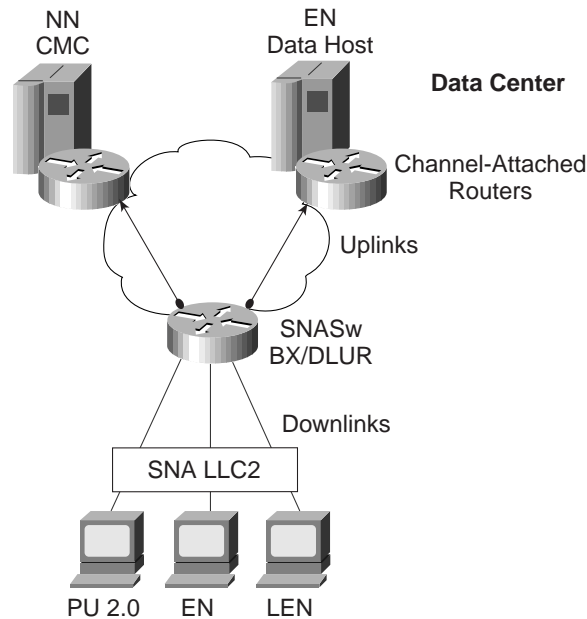
Every interface that is used for SNASw communication also requires an SNASw port definition statement. An SNASw port definition associates SNA capabilities with a specific interface that SNASw uses for transport. The port associates with a Token Ring, Ethernet, FDDI, or ATM interface that points downstream.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

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

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

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

Figure 3-3 SNASw Uplinks and Downlinks



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

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

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

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

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

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

## Supported SNASw Link Types and Transport Options

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

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

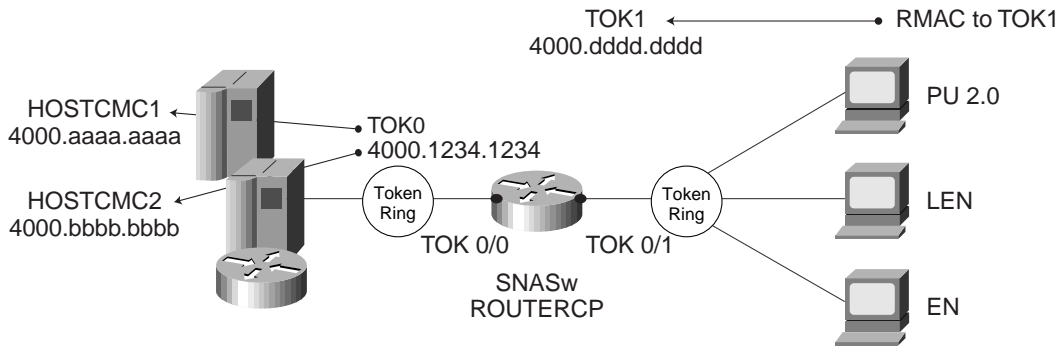
SNASw does not support the following transport options:

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

## Native LAN Transport

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

Figure 3-4 Native LAN Transport



```
interface TokenRing 0/0
mac-address 4000.1234.1234
ring-speed 16

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

**SNASw/APPN Control Point Name**  
snasw cpname NETA.ROUTERCP

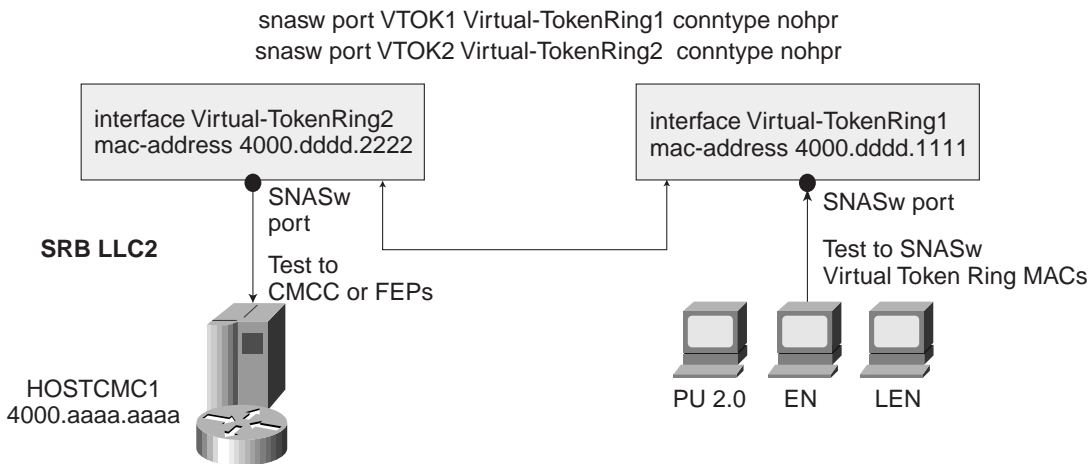
**SNASw Port for Downstream SNA Devices**  
snasw port TOK1 TokenRing0/1 conntype nohpr

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

## Virtual Token Ring Transport

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

Figure 3-5 Virtual Token Ring Transport



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

## VDLC Transport

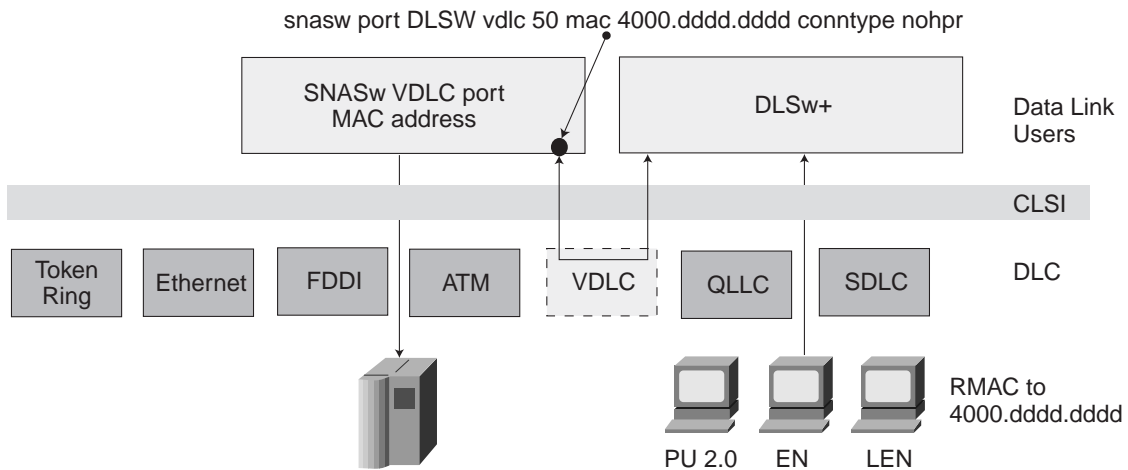
SNASw uses VDLC to connect to DLSw+ transport and local switching technologies. VDLC is used for a number of connectivity options, including the following:

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

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

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

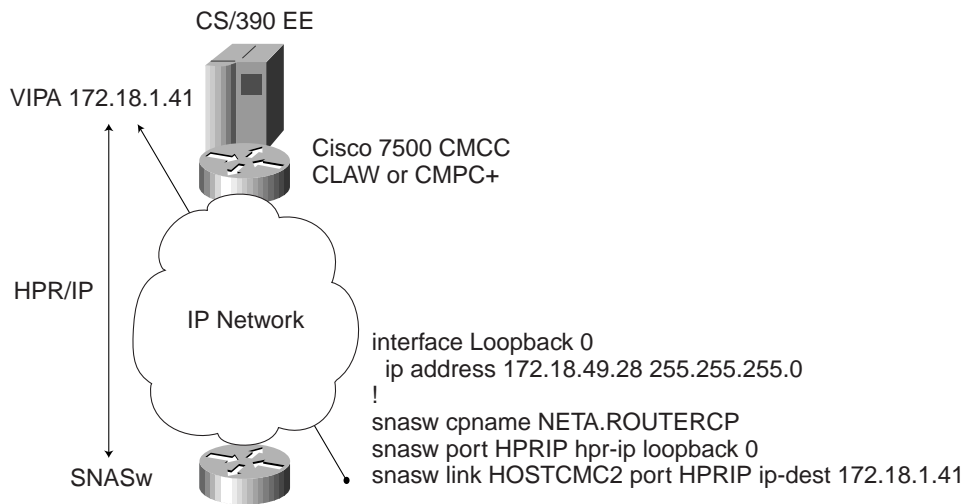
Figure 3-6 VDLC Transport



## Native IP DLC Transport

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

Figure 3-7 Native IP DLC Transport



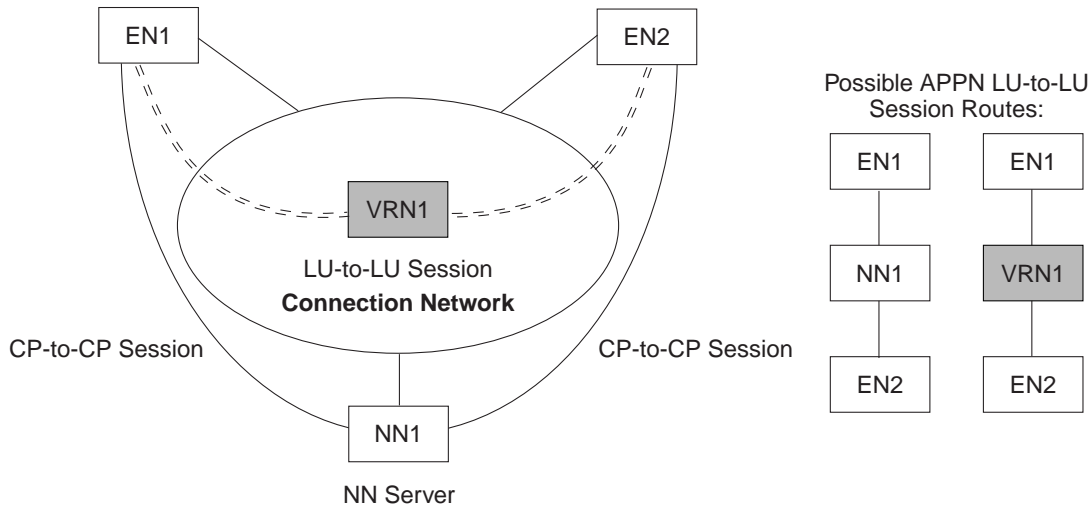
## SNASw Connection Network Support

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

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

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

Figure 3-8 VRNs and Connection Networks



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

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

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

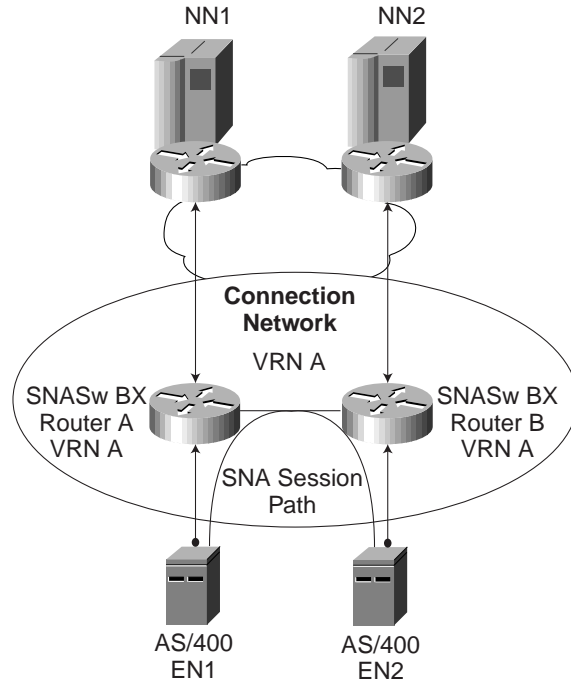
Even with OS/390 and CS/390 V2R10 (or higher) and the fix to IBM APAR OW44611 applied, there will still be one scenario that will not support HPR on an APPN link immediately adjacent to an interchange transmission group. This is when an ICN defines a connection to a connection network (VRN) and a session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

### SNASw Connection Network Support for Branch-to-Branch Traffic

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

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

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




## IBM CS/390 Connection Network Support

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

## Enabling SNASw DLUR Support

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

Figure 3-10 SNASw DLUR Support



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

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

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

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

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



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

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

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

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

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

## Customer Scenarios for SNASw Deployment

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

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

### Customers Migrating from APPN NN to SNASw

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

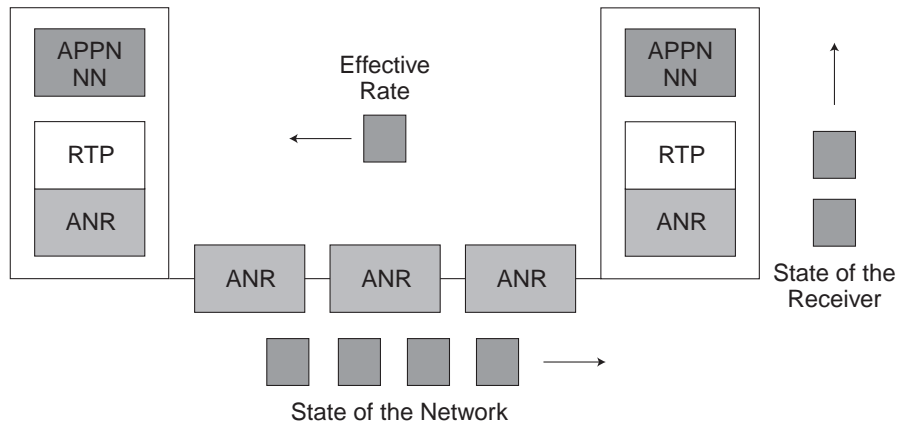
#### Memory and CPU Processing Requirements

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

#### Flow Control Issues

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


Figure 3-11 ARB Flow Control



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

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

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



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

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

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

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

## Scalability Issues

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

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

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

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

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

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

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

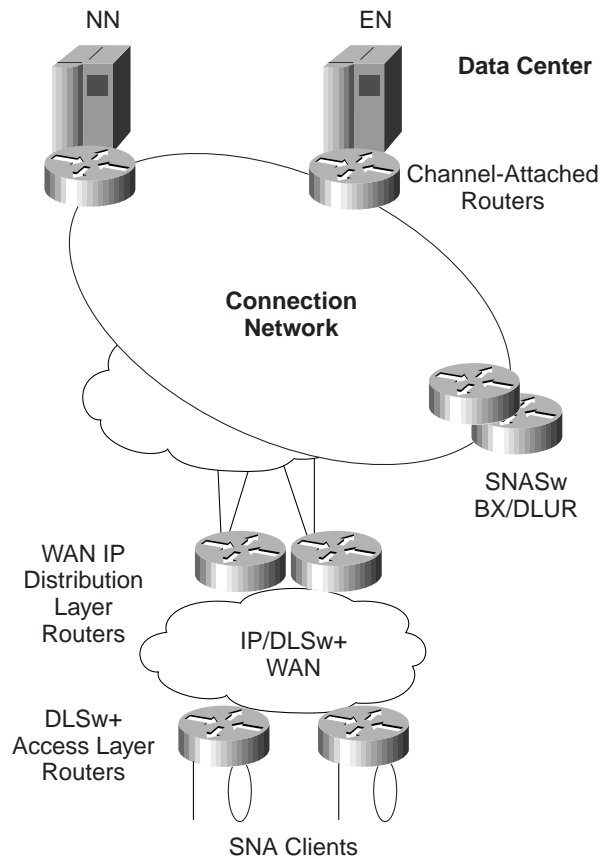
BX works with any downstream EN (except a CS/390 VTAM or PBN host—see the “SNASw General Considerations” section) to select the session path from the end user to BX. BX works with the CS/390 NN server to select a session path between BX and the application host. Links to nodes downstream of a BX are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

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

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

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

Figure 3-12 BX Network Design for Migration



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

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

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

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

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

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

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

### Leveraging SNASw EE for APPN NN-to-SNASw Migration

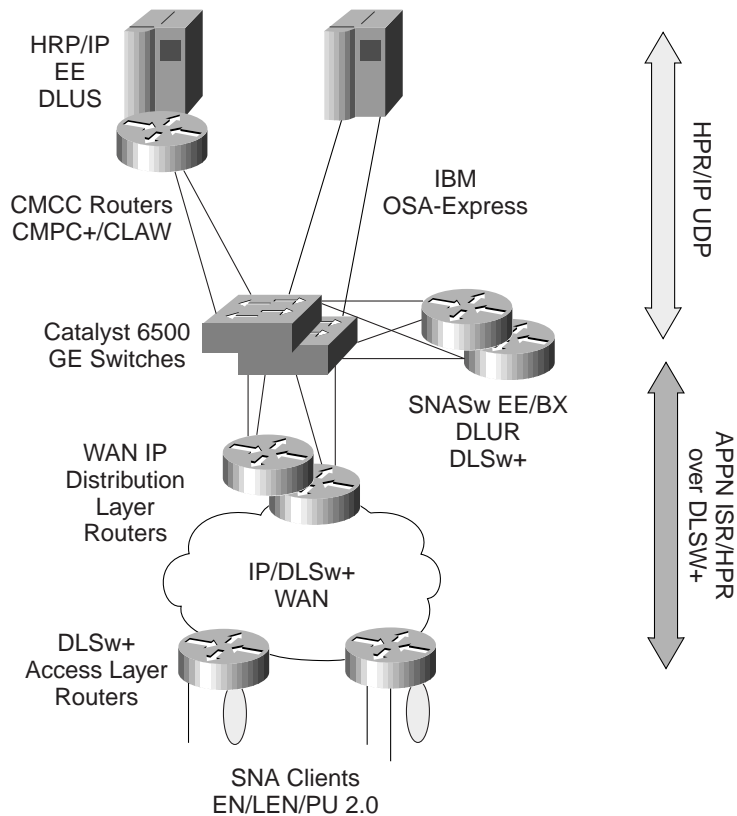
EE was created within the open framework of the AIW and submitted to the IETF as RFC 2353 in May 1998, and it has been commercially available since 1999. Cisco SNASw implements the EE feature as well as IBM CS/390 V2R6 (with APAR OW36113 applied) and higher. EE is a logical connection that represents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

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

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

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

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



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

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

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

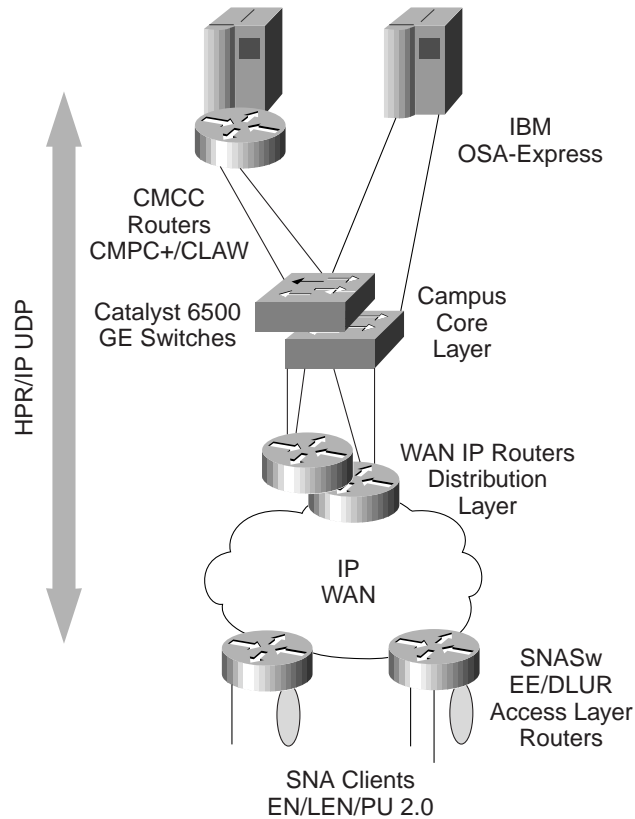
This design option illustrates the use of both Cisco channel-attached routers (CIP/CPA) for the IP uplink to the mainframe server and IBM's Gigabit Ethernet OSA-Express. Both solutions provide viable design alternatives. Typically, when the IBM S/390 (or zSeries) Parallel Sysplex is running OSA-Express,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

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

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

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

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



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

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

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

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

## Customers Migrating from Subarea SNA (FEP) to SNASw

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

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

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

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

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

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

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

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

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

## Subarea SNA to SNASw DLUR Migration Considerations

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

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

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

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

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

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

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

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

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

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

### SNASw DLUR Design Considerations

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


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

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

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

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

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



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

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

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

For a more detailed discussion and business case, see *Token Ring-to-Ethernet Migration* at [www.cisco.com/warp/public/cc/so/neso/ibso/ecampus/trh\\_bc.htm](http://www.cisco.com/warp/public/cc/so/neso/ibso/ecampus/trh_bc.htm).

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

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

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

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

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

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

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

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

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

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

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

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

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

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