

# Positioning SNASw and DLSw+

## Overview

For several years, Cisco had only one strategic technology that enabled transport of SNA over an IP backbone. That technology was DLSw+. The addition of Cisco SNASw means that Cisco provides two technologies to consolidate and transport SNA over IP.

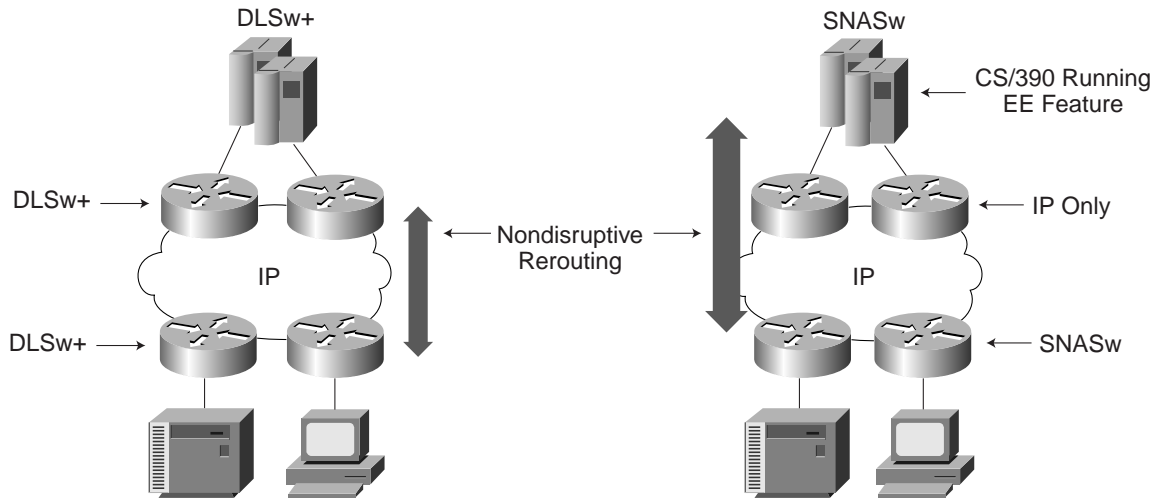
Cisco DLSw+ was created in early 1995 as a means to transport SNA over an IP network. It is based on the DLSw standards work done in the APPN Implementers Workshop (AIW) and documented in Internet Engineering Task Force (IETF) standards Request for Comment (RFC) 1795 and 2166. Cisco DLSw+ is standards compliant but includes additional features that allow it to perform better, scale larger, and provide better availability for SNA sessions.

DLSw+ enables media independence for SNA devices and simplifies SNA network configuration. It works with any SNA network: subarea, APPN, or APPN/HPR; any level of CS/390; and any LU or PU type (including PU 4). It has been a part of the Cisco IOS Software since Release 10.3 and has been continuously enhanced. To date, several hundred thousand routers run DLSw+ in production networks. Hundreds of banks, retail establishments, credit card companies, and carriers have adopted DLSw+ as their standard protocol for SNA transport over IP until now.

SNASw, Cisco's second-generation APPN platform, was created in 1999 as a replacement for Cisco's first-generation APPN NN product. It is currently available in Cisco IOS Release 12.1 and higher. APPN NN is reaching its EOE milestone concurrent with EOE for Release 11.2 on April 16, 2001, and EOE for Release 12.0 after March 2002. It is also based on standards work done in the AIW. SNASw provides two key functions: SNA routing and SNA transport in IP. SNASw uses APPN to provide SNA routing. Unlike previous APPN support in the Cisco IOS Software, SNASw includes the BX feature, an enhancement to APPN that enables it to scale. In addition, SNASw provides EE RFC 2353 support. EE enables transport of SNA data over an IP network. SNASw also provide DLUR boundary function support for dependent SNA PU 2.0 devices.

The *key* difference between the IP transport provided by SNASw EE and DLSw+ is that with SNASw EE, you can design your network to transport SNA in IP all the way into your IBM S/390 and zSeries hosts, eliminating any single points of failure in the network. SNASw EE interoperates with EE support in IBM CS/390 releases since V2R6 (with APAR OW36113) and, hence, extends the high availability characteristics of IP all the way to the EE-enabled enterprise server, as shown in Figure 2-1. With DLSw+, you can transport SNA over IP between DLSw+ peering routers, but the last hop into the mainframe is SNA.

Figure 2-1 Comparison of Availability Characteristics of DLSw+ and SNASw



The decision between these two solutions is not necessarily an either/or one. Many enterprises will deploy both. If your network currently uses FEPs for native SNA routing and you are migrating your data center from FEPs to Cisco CIP or CPA platforms (or IBM OSA-Express), then you should use SNASw somewhere in your network to provide the SNA routing function. You can deploy SNASw in the branch (as an alternative to DLSw+) or only where you currently have FEPs (in which case you can still use DLSw+ in the branches).

This chapter includes the information you need to decide which solution is appropriate for your network and where it should be deployed. It describes each technology at a high level, highlights the relevant features of each, outlines key decision criteria, describes when you need SNA routing, and suggests possible network designs.

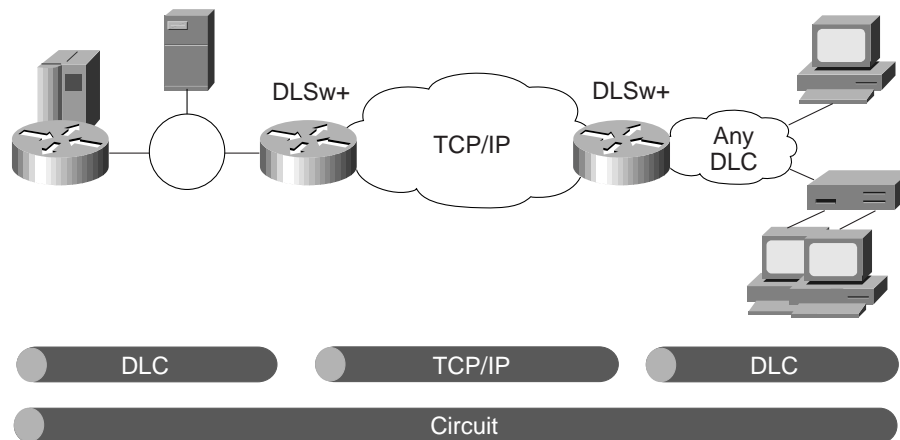
## DLSw+ Features

The DLSw+ architecture offers availability, scalability, quality of service (QoS), flexibility, and management features.

## Architecture

DLSw+ transports SNA in TCP/IP (DLSw+ also offers IP-only and direct encapsulation). To monitor and maintain an end-to-end connection, DLSw+ uses a construct known as a circuit. It is comprised of two data-link control (DLC) connections and the TCP connection, as shown in Figure 2-2. TCP provides reliable delivery of data and flow control. In addition, DLSw+ has circuit-level flow control to enable traffic from one circuit to be slowed down while traffic on other circuits is unaffected.

Figure 2-2 DLSw+ Architecture



## Availability

DLSw+ provides rerouting around link failures between the DLSw+ routers. If a link fails between the end system and a DLSw+ router, the sessions using that link will fail (unless the end systems are using HPR). Recovery from the failure, however, is dynamic.

## Scalability

There is no known restriction on how large a hierarchical DLSw+ network you can build. The number of central site DLSw+ routers can grow in proportion to the number of remote sites and the traffic volumes—ranging from one central site router for every 100 remote locations with moderate to heavy interactive traffic to one central site router for every 200 to 300 remote branches with automatic teller or point-of-sale (POS) machine traffic. The controlling factors include the number of PUs, the number of remote sites, the traffic volumes, and the broadcast replication (typically not a problem in SNA networks).<sup>1</sup>

## QoS

DLSw+ automatically sets IP precedence, ensuring that SNA traffic receives better treatment in any network that supports IP precedence. In particular, in a Cisco environment running Cisco QoS algorithms such as Weighted Fair Queuing (WFQ) and Class-Based Weighted Fair Queuing (CBWFQ), the DLSw+ traffic is placed in a queue that is serviced more frequently. In periods of severe congestion where packet loss could potentially occur, DLSw+ packets would be among the last to be dropped. Because DLSw+ uses the same flow control method as other TCP applications, networks tuned to work well with TCP flow control will work well with DLSw+.

If DLSw+ is running in a router that is also running SNASw, DLSw+ automatically maps standard APPN COS to IP precedence ToS bits in the IP header.

## Flexibility

DLSw+ is a flexible solution that addresses requirements for a variety of devices, protocols, WAN speeds, and LAN media.

1. For more information, see *DLSw+ TCP Performance* at [www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/tech/dstcp\\_wp.htm](http://www.cisco.com/warp/public/cc/pd/ibsw/ibdlsw/tech/dstcp_wp.htm).

## Device and Protocol Support

DLSw+ supports any SNA devices or SNA networking architecture. It supports subarea, APPN, and HPR. It supports communication between any PU types including PU 4 (FEPs). Hence, DLSw+ can be used between different SNA networks transporting SNA Network Interconnection (SNI) traffic.<sup>2</sup> In addition, DLSw+ supports NetBIOS traffic.

## Encapsulation Options

In addition to TCP/IP transport, DLSw+ supports UDP, Fast Sequenced Transport (FST), and direct encapsulation. FST is a high-performance encapsulation option that can be leveraged for use over higher-speed links (256 kbps or higher) when high throughput is required. FST is faster and has less overhead than TCP encapsulation, but it relies on the end systems for recovery. FST uses sequence numbers in a field in the IP header to ensure that packets arrive in order. Out-of-order packets are discarded. (DLSw+ FST transport does not support multilink transmission groups between FEPs.)

Direct encapsulation uses only link-layer encapsulation and is possible for Frame Relay or HDLC. Direct encapsulation has the least overhead but does not offer nondisruptive rerouting around link failures. In addition, it requires that the DLSw+ routers be adjacent to one another (separated only by a link or Frame Relay network), limiting network design options and forcing the DLSw+ routers to also handle WAN functionality.

## Media Support

DLSw+ supports attachment to end systems over almost all media.<sup>3</sup> However, when there are Ethernet-attached devices, network design limitations must be considered. In mixed Ethernet and Token Ring environments, the Token Ring-attached devices are limited to a frame size that will work on Ethernet. For releases prior to Cisco IOS Release 12.0(5T), in environments where Ethernet-attached devices initiate SNA connections, loops are possible. Careful design is required to prevent them. Release 12.0(5T) added a feature to enhance the redundancy characteristics when Ethernet exists in the network. See the *DLSw+ Design and Implementation Guide* at [www.cisco.com/warp/public/cc/cisco/mkt/iworks/wan/dlsw/prodlit/toc\\_rg.htm](http://www.cisco.com/warp/public/cc/cisco/mkt/iworks/wan/dlsw/prodlit/toc_rg.htm) for more information.

## Management

DLSw+ can be managed with CiscoWorks Blue Maps and SNA View. SNA View provides first-level problem isolation for DLSw+, SNASw, native SNA, TN3270, and remote source-route bridging (RSRB). Maps shows a graphical view of a DLSw+ network and provides circuit-level problem determination information. Show commands provide additional detail, utilizing an extensive DLSw+ MIB. Hop-by-hop performance can be measured with Internetwork Performance Monitor (IPM). S/390 management of Cisco routers is available with CiscoWorks Blue Internetwork Status Monitor (ISM).

## Cost

DLSw+ is a slightly less-expensive SNA transport branch solution than SNASw EE, partly because of its packaging (it is part of the IBM base feature set in Cisco IOS Software) and partly because of its size. DLSw+ branch routers require a smaller image size than SNASw routers and, hence, less memory.

2. DLSw+ transports SNI traffic between FEPs but does not replace the SNI function provided by FEPs.

3. DLSw+ has some media restrictions. For example, X.25 switched virtual circuits (SVCs) are not supported between DLSw+ and an end system. For a complete list, see the *DLSw+ Design and Implementation Guide*.

## SNASw Features


The SNASw architecture also offers SNA transport over IP, availability, scalability, QoS, flexibility, and management features.

## Architecture

The SNASw EE feature transports SNA using UDP/IP encapsulation. The RTP component of HPR provides reliable delivery of frames and flow control at HRP RTP EE endpoints (similar to TCP's function in a TCP/IP network). SNASw supports BX, which provides for greater scalability. The BX feature allows SNASw to provide emulated NN services to downstream SNA devices while providing an EN appearance to the upstream NN server in CS/390.

SNASw appears like an EN upstream (see Figure 2-3) and, hence, does not participate in topology updates and locates as APPN NNs do. This is what allows SNASw networks to scale. However, SNASw provides an NN image and SNA routing services to downstream SNA devices. It can register downstream devices to the CS/390 host NN central directory server and provide DLUR function so that non-APPN SNA-dependent devices can take full advantage of the transport availability afforded by HPR over IP (SNASw DLUR replaces the FEPs boundary function for SNA PU 2.0 dependent devices in subarea networks.)

Figure 2-3 SNASw Architecture



## Availability

SNASw supports HPR, which provides nondisruptive rerouting around link failures between HPR RTP endpoints. Figure 2-3 shows a CS/390 enterprise server at one end and a Cisco router running SNASw at the other end. HPR could also be running in two CS/390 enterprise servers, or one CS/390 host and an EN that supports HPR.

The most common reason to implement SNASw and HPR support is to enable a router to communicate directly to IBM S/390 or zSeries hosts, encapsulating SNA in IP and enabling nondisruptive rerouting around channel or data center router outages.

## Scalability

The key scalability question is how large a network can one support using SNASw. This depends upon where you implement SNASw, either in a branch servicing the resources for the branch, or in a distribution site or data center servicing the resources for multiple branches.

If SNASw is running in a branch, the limiting factor is not SNASw, but CS/390. Joint IBM and Cisco testing has demonstrated that the EE function in CS/390 scales to support at least a 2000-branch network. In these tests, SNASw EE was running in multiple routers and communicating with a single logical partition (LPAR) on the

S/390 host. Ten thousand RTP connections were established (each remote site requires multiple RTP connections, including control sessions and one RTP per COS). Thirty thousand LU-to-LU sessions across these ten thousand RTP connections were set up and maintained without any noticeable response degradation, with less than 30 percent of the S/390 CPU being utilized.

If SNASw is not running in a branch, then the scalability question is different. When SNASw runs in a distribution site or in the data center, it minimizes the number of RTP connections to CS/390. Hence CS/390 scaling is not an issue. In a distribution site or data center, SNASw also provides DLUR and NN services to the branches. The key scalability question is how many downstream resources can a single SNASw router support. To support large numbers of downstream SNA resources, multiple SNASw routers may be required. Like DLSw+, SNASw should scale to any size network. The number of central site routers will grow proportionately to the number of remote sites. The controlling factors are the number of LUs and the traffic volumes.

## QoS

SNASw automatically maps standard APPN COS to IP precedence, ensuring that SNA traffic receives better treatment in any network that supports IP precedence. In particular, in a Cisco environment running Cisco IOS QoS capabilities such as CBWFQ, the high-priority SNA traffic is placed in a queue that is serviced more frequently. Low-priority SNA packets are placed in a queue that is serviced less frequently, but more frequently than a queue with unmarked packets. In periods of severe congestion, high-priority SNA packets would be among the last to be dropped.

## Flexibility

SNASw has several characteristics that should be considered when assessing the appropriate implementation of the two options (DLSw+ and SNASw).

### Device and Protocol Support

SNASw supports SNA only. SNASw BX support requires CS/390 to be at V2R5 or higher in your S/390 or zSeries host, and CS/390 must be running APPN. To take advantage of the SNASw EE feature (to enable SNA transport over native IP all the way into the host), you must be running CS/390 V2R6 (with APAR OW36113 applied) or higher, and CS/390 must be running APPN/HPR. Because SNASw includes DLUR support, it supports communication between PU 2.0 dependent SNA devices and the CS/390 SSCP.

### Encapsulation Options

SNASw with the EE feature uses UDP encapsulated IP as the SNA transport mechanism from the SNASw EE router to the EE-enabled enterprise host server. RTP provides reliable delivery and flow control utilizing Responsive Mode adaptive rate based flow control. From the perspective of SNA APPN, EE is just another DLC connection type; to the IP network, EE is just another UDP application running in your network!

### Media Support

SNASw has no media restrictions. When running SNASw in an Ethernet environment, EE provides redundancy capability and dynamic routing capabilities inherent in IP networks. Loops are not an issue as they are with DLSw+ Ethernet environments, because SNA transport between the SNASw EE router and the EE enterprise server is entirely over Layer 3 IP.



## Management

SNASw can be managed with CiscoWorks Blue Maps and SNA View. SNA View provides first-level problem isolation for DLSw+, SNASw, native SNA, TN3270, and RSRB. Maps shows a graphical view of a Cisco SNASw network and provides session-level problem determination information. Show commands provide additional detail, utilizing an extensive SNASw MIB. Hop-by-hop performance can be measured with IPM. S/390 management of Cisco routers is available with ISM.

When SNASw EE is used end to end throughout the network, SNA traffic is essentially UDP over IP application data (the entire network is IP). Network management can therefore be performed using the Cisco IP network management tools (CiscoView 2000).

## Cost

Routers running SNASw cost slightly more than routers running DLSw+ because SNASw is a separate Cisco IOS feature not included in the base IBM feature set. However, it should be noted that running SNASw EE out to remote branch locations does *not* require an SNA router in the data center, because SNASw EE transport is native IP from the remote branch SNASw router into the EE-enabled S/390 or zSeries enterprise host, thus reducing the overall cost of the SNASw EE transport solution.

## Decision Criteria

For many customer environments, deciding between these two technologies is very straightforward and in some cases it is not. This section leads you through the questions you need to answer to determine if one or both of these technologies combined are appropriate in your network.

### Is SNA Routing Required?

If the target SNA application is not in the same SSCP domain as the SSCP-to-PU and SSCP-to-LU control sessions, SNA routing is required. Without SNASw or APPN, this routing is done by the owning SSCP itself, forcing the data center hosts to perform SNA routing and raising the application cost by increasing mainframe CPU utilization. Therefore, if SNA routing is required, you need SNASw. SNA routing, which is required in any cross-domain environment, has traditionally been done by FEPs, but the majority of customer enterprises today have chosen to migrate from FEPs to higher-speed, multiprotocol routers to save money and to position their networks for multiprotocol data/voice/video applications supported by the Cisco Architecture for Voice, Video, and Integrated Data (AVVID), as well as applications such as IBM WebSphere on S/390 and zSeries Parallel Sysplex mainframes. If you want to replace your FEPs with routers, and your FEPs are currently routing SNA traffic between different LPARs, then you will most likely want your Cisco routers to provide that same functionality. SNASw provides that capability. You may still need to decide where you want to run SNASw—that is, at the branch or at a distribution site. That decision is covered in the “Network Design” section.

### Are You Ready?

If you want to consolidate SNA and IP traffic onto a single backbone, then you can use either DLSw+, SNASw, or DLSw+ and SNASw combined. Before you weigh the technical merits of each, answer these simple questions:

- *What operating system level are you running on your S/390 or mainframe?*—The SNASw BX feature requires CS/390 V2R5 or higher. The EE feature in CS/390 requires OS/390 V2R6 (with APAR OW44611 applied) or higher.
- *What Cisco IOS release are you running on your branch routers?*—SNASw requires Cisco IOS Release 12.1 or higher.
- *Are you running APPN or APPN/HPR in CS/390 today?*—The biggest advantages SNASw brings is the ability to provide necessary SNA routing (BX feature) and transport SNA in native IP all the way into your S/390 (EE feature). However, to take advantage of those capabilities, you need to be CS/390 APPN-enabled (for BX) and APPN/HPR-enabled (for EE) and be at the required version and release of CS/390 software.

Additional considerations include availability, cost, and application direction. SNASw offers the highest potential availability (although it requires much newer S/390 and router software). It costs slightly more per branch router than DLSw+ because of the additional cost of the SNASw Cisco IOS feature set. If you intend to continue supporting existing mission-critical SNA applications on your S/390 and zSeries mainframes for several more years while developing newer applications in IP, you should consider SNASw because it can bring you the highest availability possible. You may want to start with SNASw only in the data center (combined with DLSw+ in data center peering routers), and then migrate SNASw EE out to distribution sites or remote branch offices. See the “Network Design” section for more detail.

## Network Design

The technologies described in this chapter are not mutually exclusive. Both SNASw and DLSw+ can be used in the same network, and in fact, a large number of Cisco customer implementations of have done just that.

In many customer situations, DLSw+ supports SNA WAN transport to remote DLSw+ peering routers at remote branch locations. SNASw can be deployed in locations where the network has FEPs to replace necessary SNA routing previously provided by the FEPs, and to replace FEP boundary function support for dependent SNA devices. If FEPs are in the data center or distribution sites, that is where the SNASw routers reside. At the data center or distribution site, SNASw can run in the same routers that currently handle DLSw+. This implementation allows organizations to isolate the SNASw function to an “SNA router”—so named because it provides the functions required to support necessary SNA routing of client sessions directly to the target application host in addition to providing DLSw+ peer termination points for WAN transport of SNA from remote SNA devices.


However, there are many network design possibilities. Rather than go through them all, this chapter discusses three hypothetical case studies, explains the chosen design, and describes the reason for the choices. Because the focus of this chapter is to describe DLSw+ and SNASw as SNA-over-IP transport alternatives, none of these case studies shows SNASw being deployed without the SNASw EE feature (although it is certainly possible to do so predicated on the environment and requirements).

### Case Study 1—No HPR over IP Transport (DLSw+ Only)

In this case study, the enterprise has no requirement for APPN because it has no cross-domain sessions, is not using FEPs to route cross-domain traffic, or routes cross-domain traffic through CS/390 enterprise servers. The enterprise is not planning to change its SNA applications and is planning to maintain the status quo in the data center for some time under the reasoning “it works; therefore, don’t touch!”

The simplest, surest, and lowest-cost solution for this customer is to use DLSw+ to transport SNA traffic over an IP backbone back to a data center router. The customer chose to keep the SNA data center router separate from the WAN distribution router to simplify change management and maximize availability. The SNA data center router runs the Cisco IOS level that has all the DLSw+ features needed, and the customer does not want to modify it to pick up the latest compression or security feature (and vice versa for the WAN distribution router). Two CIPs (one primary and one backup) run Cisco SNA (CSNA) to handle all the SNA traffic, and four Cisco 7200 Series routers run DLSw+ (to handle a 600-branch network), including one DLSw+ router used only for backup. Figure 2-4 shows this design.

Figure 2-4 DLSw+ Design



## Case Study 2—DLSw+ and SNASw


In this case study, the enterprise wants to leverage its Parallel Sysplex complex and achieve the high availability it affords. The customer is migrating to Cisco Catalyst® 6500 Gigabit Ethernet switches with Multilayer Switch Feature Cards (MSFC) installed for IP Layer 3 support in the data center, and new host applications are being written to run IP natively. However, it will be several years before a complete migration from SNA to IP applications is complete, and in the interim, the customer wants the high availability and design simplicity afforded by having an all-IP data center today.

This enterprise already uses DLSw+ to transport SNA traffic over an IP backbone. The customer chose not to deploy SNASw EE out to the remote branches at this point because the DLSw+ network has been in place and quite stable for a long time (if network outages occur over the DLSw+ network, they affect only a small portion of the network and are recovered automatically). However, the customer wants to ensure that a CIP or channel outage (which today would bring down almost the entire network) can be handled transparently and nondisruptively. Hence, the customer is adding SNASw to hub-end data center routers terminating DLSw+ peer connections coming in from remote branch DLSw+ routers. The BX capability of SNASw is providing necessary SNA routing for downstream SNA devices, while at the same time the SNASw EE feature transports SNA traffic natively over IP into the CS/390 enterprise server upstream. By doing this, the customer has eliminated Token Ring source-route bridge requirements for maintaining multiple active redundant bridged routing information field (RIF) paths to the mainframe (required for Logical Link Control, type 2 [LLC2] bridged transport), and has also eliminated the potential for loop problems that can occur in a bridged Ethernet environment. In addition, should a channel failure occur, IP immediately reroutes traffic and SNA sessions are not impacted (since EE uses HPR for nondisruptive session path switching). Finally, this design positions the customer to use Cisco Catalyst 6500 Gigabit Ethernet switches connected to IBM S/390 OSA-Express for SNA traffic transport over IP (HPR/IP).<sup>4</sup>

4. The OSA-Express Gigabit Ethernet card is for TCP/IP environments only. This card only supports SNA traffic when SNA is encapsulated in IP using the EE support in OS/390 Version 2 Release 6 (with APAR OW44611) or higher.

As in the previous case study, this enterprise chose to keep the SNA data center router separate from the WAN distribution router to simplify change management and maximize availability. Two CIPs (one primary and one backup) run IP to handle all the SNA traffic, and six Cisco 7200 Series routers run DLSw+ and SNASw (to handle a 1000-branch network), including one DLSw+/SNASw router used only for backup. Figure 2-5 shows this design.

Figure 2-5 Combined SNASw and DLSw+ Design




### Case Study 3—SNASw Only

In this case study, the enterprise demands the highest availability for its SNA applications. The customer has invested a great deal in rewriting applications to LU 6.2 and wants to continue to leverage that SNA application investment. The customer was running separate networks for SNA and IP and decided to consolidate using SNASw EE HPR over IP support for SNA transport natively over IP. The network is already at the latest operating system level and is running APPN/HPR and EE support in CS/390.

The customer has 200 regional offices that run SNASw EE. From the branch into the S/390 host, the SNA traffic is transported in IP. Hence, there is no need for SNA routers in the data center. The customer leverages the Cisco QoS features to ensure that the interactive SNA and Telnet traffic take precedence over SNA batch and FTP traffic. Figure 2-6 shows this design.

Figure 2-6 SNASw Design



## Conclusion

Cisco DLSw+ is a proven method of transporting SNA traffic over an IP network. The downside of DLSw+ is that the final hop into the mainframe is SNA. SNASw EE can eliminate this issue by combining the best qualities of APPN and DLSw+. SNASw is going to play a key role in many data centers in the future. Whether it is right for you today depends on your application direction, cost constraints, OS/390 level, and CS/390 configuration. If you are moving to APPN and APPN/HPR, SNASw can provide enhanced availability and data center flexibility. DLSw+ will also continue to play a key role as a WAN transport for SNA traffic for the foreseeable future.

