

SNASw Enterprise Extender—Building a Modern IP Network While Continuing to Support SNA Devices and Applications

The phenomenal success of the Internet and World Wide Web have driven most organizations to add support for networking infrastructure and applications based on Transport Control Protocol/Internet Protocol (TCP/IP). Nonetheless, many organizations continue to utilize applications, devices, and networking infrastructure based on IBM's Systems Network Architecture (SNA) or its follow-on, Advanced Peer-to-Peer Networking (APPN). The reason? SNA has formed the basis of mission-critical systems and applications that have been developed, enhanced, and tested over a period of decades. SNA is often the underlying architecture of the applications that are at the very heart of the IT infrastructure, including customer databases, customer service transactions, financial records and applications, manufacturing resource planning (MRP) systems, and newer enterprise resource planning (ERP) systems, to name a few. Some organizations have even made information from SNA applications available to employees, customers, and partners through a modern, graphical Web interface.

Organizations that had both SNA and IP applications and infrastructure have had two basic choices in the past. They could keep the two networks separate, or they could migrate to a common IP backbone and begin to integrate the two environments together. Although it is expensive to maintain two different networks, some organizations accepted the cost in order to minimize the disruption to the mission-critical SNA applications.

However, the era in which maintaining two separate networks is a viable option is nearing its end. The IBM 3745 and 3746 Controllers, which form the backbone of an SNA networking infrastructure, have been withdrawn from marketing effective September 27, 2002. After that date, only used products and features will be available. Networks that are based on these controllers will begin to atrophy; eventually the controllers will no longer be supported. The reason the controllers have been withdrawn from marketing is the success of IP in corporate networks. "The explosive growth of the Internet and TCP/IP traffic have resulted in a severe decline in the demand for new 3745 and 3746 Communication Controllers. Thus, the products are being withdrawn from marketing..." (IBM Announcement Letter; February 26, 2002). These controllers run Advanced Communications Function/Network Control Program (ACF/NCP) software and are referred to as front-end processors (FEPs).

Although the two basic choices have been to maintain separate networks or to form an IP backbone, in reality many organizations are somewhere between these two extremes. They have a variety of technologies employed, including perhaps traditional SNA, APPN, IP, and Data Link Switching (DLSw). The withdrawal of the IBM networking controllers provides an impetus for all organizations with existing SNA infrastructure to evaluate their networks and to devise a plan to migrate to a new network.



This new network must be based on IP, today's networking standard, but it must also seamlessly support all of the organization's SNA and APPN applications and devices.

There is no longer any doubt that IP is *the* strategic networking protocol for the future. It has been around since 1969, at the inception of the early version of today's Internet. It has proven to offer scalability, redundancy, resiliency, openness, adaptability, and manageability. The incredible popularity and exponential growth of the Internet and the Web have ensured its choice as the networking architecture for the early part of the 21st Century. Because of this, the networking industry is working, as a whole, to continue to enhance IP and its related technologies. The same cannot be said of SNA. By pushing SNA and APPN to the very edges of the network (that is, the device and application), the entire network can benefit from the continual enhancement in IP, all the while protecting the investment in SNA and APPN applications.

Two alternatives exist: SNA Switching Services (SNASw) and Cisco's Data Link Switching Plus (DLSw+). For many organizations that still utilize FEPs, the optimal solution is SNASw using Enterprise Extender (EE). EE offers organizations the ability to keep existing SNA devices and applications, but support pure IP from end to end with no loss of availability or reliability.

An additional capability of FEPs is SNA Network Interconnect (SNI), the ability to connect different SNA networks. This is very often used to securely connect different organizations for transactional or batch exchanges and is, for some organizations, the only remaining function being used on their FEPs. EE, in combination with host-based Border Node capability, provides a complete functional replacement for SNI. Because this solution creates a totally IP-based network, from host to host, there is no need for SNASw, or any other SNA-related function, in the network. For this reason, this particular application of EE will not be discussed further in this paper.

The purpose of this paper is to describe the modern IP network, evaluate the potential solutions based on environment, and provide a road map for action.


Building the Modern IP Network

Organizations that support both SNA and TCP/IP in their enterprise networks have had two basic choices: keep the networks separate, or integrate them at some level. Within those two basic choices there is an infinite number of potential combinations of products and technologies. However, in general there are four general design options: physically separate, physically integrated, logically integrated backbone, and logically integrated end to end.

Figure 1 depicts the first basic design option. In this design, the SNA network and the TCP/IP network are completely separate. Each network contains dedicated equipment and communication lines. The S/390 Enterprise Server may actually participate in both networks, because TCP/IP on the mainframe has been commonplace for many years.




Figure 1: Physically Separate Network Design



Obviously, this design can be very expensive to establish and operate. There is considerable redundancy and waste due to the duplication of communication lines, networking hardware equipment, maintenance expenses, software license fees, management tools, and staff to install, monitor, and manage the two environments.

The first step in eliminating some of this redundancy is to physically provide a common infrastructure that is used by both SNA and TCP/IP. Figure 2 depicts this design. Here, both protocols share the same LANs, communication lines, and certain networking devices, particularly bridging routers. This level of integration is still only physical, not logical, because the traffic from the two protocols are like ships in the night—they are physically close, but there is no interaction.

Figure 2: Physically Integrated Network Design



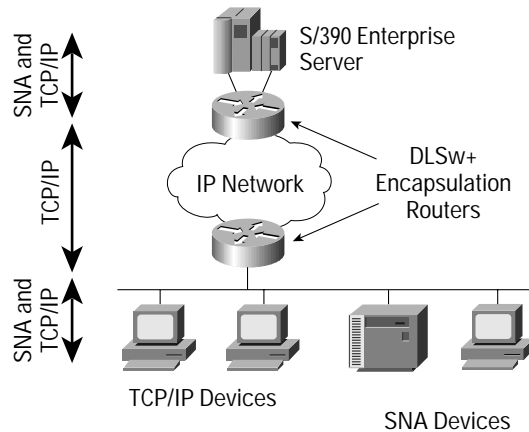
Although this design eliminates some of the redundancies and added cost inherent in the previous design, it is still very costly to configure, administer, and manage the physically integrated network.

For this reason, thousands of organizations worldwide have taken the step of logically integrating the backbone of the SNA and TCP/IP network. Cisco's DLSw+ is the leading solution for providing this integration and has been deployed on several hundred thousand routers around the world. Figure 3 depicts a logically integrated backbone



network design achieved by implementing DLSw+. SNA traffic is encapsulated in IP at the edges of the network. The backbone itself is pure IP. The SNA data gains all of the benefits inherent in an IP network, including the ability to design redundant paths and components and the automatic routing around failed components and links.


Figure 3: Logically Integrated Backbone Network Design



This design is a vast improvement over the previous designs and has been proven in countless networks, both large and small. However, because the DLSw+ approach encapsulates one protocol (SNA) within another (IP), there is some loss of native SNA functionality. And, when a FEP is removed from the network, some networking functionality must be taken on by ACF/Virtual Telecommunications Access Method (VTAM), causing increased utilization of S/390 host resources. These issues will be discussed in some detail in the following sections.

The final approach to integration is to have an integrated end-to-end solution, in which SNA and APPN applications continue to have full functionality but utilize IP as the networking transport. With this approach, there is no loss of SNA routing and prioritization functionality, and existing SNA devices continue to be supported. Figure 4 depicts this network design.

Figure 4: Logically Integrated End-to-End Network Design





The end-to-end design depicted in Figure 4 is optimal because it achieves two major goals simultaneously: it preserves the vast investment in SNA and APPN devices and applications, and it allows them to take advantage of all of the advancements that have been made and will continue to be made in IP.

The issue, of course, is how an IT organization is going to migrate along the path to the eventual achievement of the goals depicted in Figure 4. Does it have to be a long, stepwise transition? Does this entire network have to change overnight? Will the migration cause many intermediate steps that eventually will need to be abandoned? The answer to all of these questions is “No.” Of course, the best plan and solution depends on the individual nature and needs of each network. But in many or all cases, proper planning can allow IT organizations to implement a smooth transition that minimizes risk and disruption and to minimize any investments in “throw-away” transitional equipment.

Many organizations have undertaken the migration to a logically integrated end-to-end network design. One example is Informations-Technologie Austria GmbH (iT-AUSTRIA), which migrated from a design similar to the physically integrated network shown in Figure 2. Located in Vienna, iT-AUSTRIA is the data processing outsourcing subsidiary of two major financial groups: the Bank Austria/Creditanstalt Group and Erste Bank & the Austrian Savings Banks. With more than 600 employees and a turnover of EUR 262 million last year, iT-AUSTRIA is Austria’s biggest data processing center, serving the country’s largest banking institutions. It maintains and operates the systems that deliver around-the-clock online services to locations all over Austria and internationally.

Along with common computing services such as system management and production control, iT-AUSTRIA also provides planning, management and support of the LAN and WAN infrastructure, extensive banking telephony, and growing internal Internet Service Provider (ISP) operations for their client banks. Computer operations are distributed across five separate production data centers situated some 10 km apart from one another within the confines of Vienna, with more than 7000 km of dedicated dark fiber cable connecting the sites. More than 15,000 MIPS of S/390 CPU power is required to process the 20 million online banking transactions each day for their clients’ businesses.

iT-AUSTRIA has undergone a drastic transition in a short period of time. At the beginning of the transition, the company’s network most closely resembled Figure 2, using Frame Relay as transport layer. Now more than 1000 Cisco routers are running SNASw EE. The eventual goal is to have a network that more closely resembles Figure 4. The dramatic feat has been accomplished in only six months, and the migration has been largely without problems. Josef Killmeyer, Team Leader in charge of network design, implementation, and support for iT-AUSTRIA, says: “SNASw EE is delivering everything it promised, and I am very happy that we have it in production! Our systems are ready to fallback anytime and easily, regardless of whether an outage was planned or not.”

Although this is an impressive feat, iT-AUSTRIA is not alone. Thousands of organizations have achieved logical integration in the backbone using DLSw+. Many others have successfully completed migration projects using SNASw technology. Still, some IT professionals are reluctant. They may have heard of negative effects of migration. They may have been told that certain things might “break” if migration is undertaken.

In large part, these perceptions exist because the integration technology has been changing and evolving over time. Vast improvements have been made since the early products, which debuted in 1995 with Cisco’s initial release of DLSw+. For example, SNASw EE is built on APPN technologies. Early APPN implementations had issues, particularly with scalability, but these have all been addressed in the new products and technologies. Table 1 summarizes some of the perceived barriers of integrating SNA and IP and today’s reality. Some of these perceptions will be discussed in more depth within this paper.



Table 1 Perceived Barriers of Integration vs. Today's Realities

Perceived Barrier	Today's Reality
Integrated networks will not scale to large sizes	Many Cisco DLSw+ customers support very large networks; the new SNASw technologies do not suffer from the same scalability issues that early APPN did and can also support very large networks
An integrated network is not manageable	There are many tools to manage IP networks, and Cisco offers value-added management solutions for managing SNA elements over an IP network
SNA sessions may time out if the IP network has congestion	Both DLSw+ and SNASw EE offer new traffic prioritization methods that allow SNA traffic priority to be preserved, minimizing the potential for loss of sessions; also, IP networks are inherently failure-resistant
Not all device types can be supported	Any traditional SNA device type, including logical unit (LU) 0 and LU 6.2, can be supported by SNASw EE; in addition, DLSw+ supports NetBIOS and physical unit (PU) 4 traffic
The mainframe will become overburdened with protocol processing if FEPs are removed	This is not true with SNASw EE—Cisco and IBM have jointly tested 30,000 sessions to a single logical partition (LPAR) with satisfactory mainframe utilization
Routers are a single point of failure for SNA sessions, but FEPs can be configured to provide redundant paths	With DLSw+, high availability designs can minimize the likelihood of a single point of failure; with SNASw EE, there is no single point of failure
The mainframe operating system needs to be at the latest and greatest version	The operating system level depends on the solution; SNASw EE requires that it be newer than about 1999 and enabled for APPN, but DLSw+ does not

Table 1 lists only some of the most common perceived barriers to integrating SNA and IP. Some of these barriers were once real, but advancements in technology have eliminated them. For further information on any of the barriers and today's reality, contact your Cisco sales representative or authorized partner.

Integration Technologies


To understand the benefits of one approach over another, it is first necessary to understand a little technical detail of the solutions. This section describes DLSw+, APPN, and SNASw at a high level.

DLSw+

DLSw+ has evolved since its introduction in 1995, and it offers a very rich set of capabilities and options. However, its basic function has not changed. In essence, DLSw+ provides a transport for SNA traffic in an IP network. It does this by encapsulating SNA within IP. Figure 5 depicts a common DLSw+ design, in which DLSw+-capable routers are implemented in each branch and also in the data center. Traffic within the branch remains SNA until it must traverse the backbone, at which point the DLSw+ router encapsulates it in IP. When this branch traffic reaches the data center over the IP backbone, the DLSw+ routers in the data center unencapsulate the SNA data and send it to the host, either through a FEP or by directly connecting to the host via a Cisco direct channel connection.



Figure 5: Typical DLSw+ Design



DLSw+ offers many benefits:

- Stable and rich solution that has been enhanced over the years and deployed in hundreds of thousands of routers
- High availability due to the inherent ability of IP to reroute around failed components
- Based on industry standards developed by the APPN Implementers Workshop (AIW) and documented in Internet Engineering Task Force (IETF) informational Requests for Comments (RFCs)
- Prioritization of SNA traffic to minimize the possibility of dropped SNA sessions
- Flexible solution that supports a wide variety of devices, protocols, WAN speeds, LAN media, and encapsulation options
- Support for traffic from all SNA PU and LU types (including FEPs—PU 4)
- Support for NetBIOS devices

The major downside to DLSw+ is related to the fact that it is a transport solution, not a routing solution. DLSw+ is an IP transport solution for SNA/NetBIOS between peering routers. In traditional SNA networks, ACF/VTAM software on the host and ACF/NCP software on the FEPs together perform all routing of SNA traffic, ensuring that the traffic flows appropriately from source to destination. Because it does not have any ACF/NCP or ACF/VTAM functionality within it, DLSw+ routers cannot route traditional SNA traffic. Instead, DLSw+ simply encapsulates the traffic in IP. Within the IP network the traffic is routed, but when converted back to SNA, it must continue to be routed by ACF/NCP or ACF/VTAM.

Two issues that result from this fact have caused some IT organizations to continue to support FEPs in the data center. First, many SNA networks implement a design in which two FEPs running ACF/NCP can act as backup for one another. If one of the paths to one of the FEPs fails, the SNA traffic automatically can be sent through the other FEP. Known as duplicate Token Ring interface coupler (TIC) addressing, this feature has been the centerpiece of building high-availability SNA networks. DLSw+, although offering many high-availability features, cannot directly replace this function. Therefore, SNA sessions can fail if one of the DLSw+ routers fails. Some IT organizations have continued to support FEPs in the data center primarily to provide this redundancy for SNA sessions.

The second issue is host cycle conservation. As stated, ACF/VTAM and ACF/NCP work together to route SNA traffic. If ACF/NCP is not present or is eliminated, then ACF/VTAM must take responsibility for all of the SNA routing tasks. For organizations that do not have ample host CPU cycles idle, the increase in host processing can have an adverse impact on application response and network performance.

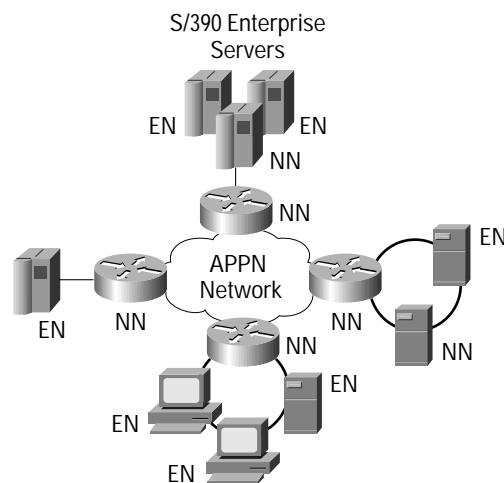


Despite these issues, many organizations have successfully implemented large-scale DLSw+ networks and have maintained high availability and conservation of host resources. DLSw+ supports the logical integration of networks and is a viable option for many environments. When coupled with SNASw EE, all of the negative issues are eliminated and organizations can confidently begin planning to eliminate their FEPs.

APPN

When IBM first introduced APPN, it was positioned as the follow-on to subarea SNA and essentially an alternative to IP networking. The original APPN architecture defined three different node types: low-entry networking node (LEN), end node (EN), and network node (NN). The backbone of this first-generation APPN network was built of NNs that provided networking services to LENs and ENs. In particular, NNs provide the native SNA routing functionality that had been performed by ACF/NCP and ACF/VTAM in a traditional, subarea network. Figure 6 depicts an early APPN network.

Figure 6: First-Generation APPN Network



The early APPN architecture had some serious limitations that impeded its widespread use. One of these early limitations was that SNA sessions would be lost if a link between two SNA endpoints failed. The members of the AIW worked together to address this and developed an enhancement to APPN called High Performance Routing (HPR). HPR provides dynamic routing around failed components, which has always been a hallmark of TCP/IP networks.

HPR added significant value to APPN, but there were still three major issues that prevented widespread adoption. First and foremost, the popularity of TCP/IP, especially after Web technologies became so pervasive, made IT executives wonder why it made sense to maintain two network protocols, one based on a new SNA and one based on TCP/IP. By the late 1990s, the majority of enterprise network executives had decided that TCP/IP would be the strategic networking protocol. This critically weakened any support for networks built using native APPN protocols.



The second major issue is that the design of APPN prevented the ability to build large networks. The main reason is that native APPN results in a large volume of broadcast traffic, because nodes communicate to locate resources. Broadcast storms, in which the network is deluged with messages in a given period of time, were common and caused instability in the network. The problem grew worse as the network size grew, so networks were effectively limited to a few hundred NNs at most—a size not sufficient to support very large enterprises.

The third major issue of native APPN is that its implementation is complex. It is difficult to configure and implement APPN nodes. When implemented, it is difficult to manage and to troubleshoot the network. And because of the lack of commercial success of APPN, there are few tools and skilled talent available to implement and run the networks. Early implementations of APPN networks revealed these issues. Customers responded with clear feedback that they wanted changes to APPN. They wanted a solution that would provide native SNA support in a manner that would provide high availability, but that would also integrate with IP. They wanted a solution that would scale to support large networks without excessive broadcast traffic. They wanted a solution that was less complex to implement and manage.

The members of the AIW listened to this early APPN customer feedback and, in 1998, responded with two solutions called Branch Extender (BX) and EE, which are both included in Cisco SNASw. These new solutions use some of the early APPN technology, but do so in a way that completely addresses the customer feedback and the lessons learned from early APPN implementations. As a result, Cisco has completely replaced its early APPN NN solution with the new SNASw.


SNASw

There are two solutions in SNASw, each designed to address different network environments and requirements. The BX solution is designed for networks that have already implemented the early APPN technologies. BX solves the scalability problems of early APPN networks by appearing to be an EN to the host while performing NN functions for downstream ENs and LEN nodes. This node emulation reduces the number of actual NNs, thereby reducing both the network broadcast traffic and the complexity. For more information on BX, see http://www.cisco.com/warp/customer/cc/pd/ibsw/snasw/tech/snst_rg.htm.

SNASw EE is also sometimes called HPR/IP (or HPR over IP). As that name implies, EE is a solution that allows APPN HPR (and traditional SNA) traffic to be transported natively over an IP network. With EE, the SNA/APPN elements are at the very fringes of the network. The applications and devices that have supported mission-critical operations for decades can continue to function, but the entire network is based on IP. As advances in IP-related technologies become available, the entire network can participate and benefit from the enhancements. The logically integrated end-to-end network described as a goal earlier in this document is achieved with EE. Figure 7 depicts a network using EE.



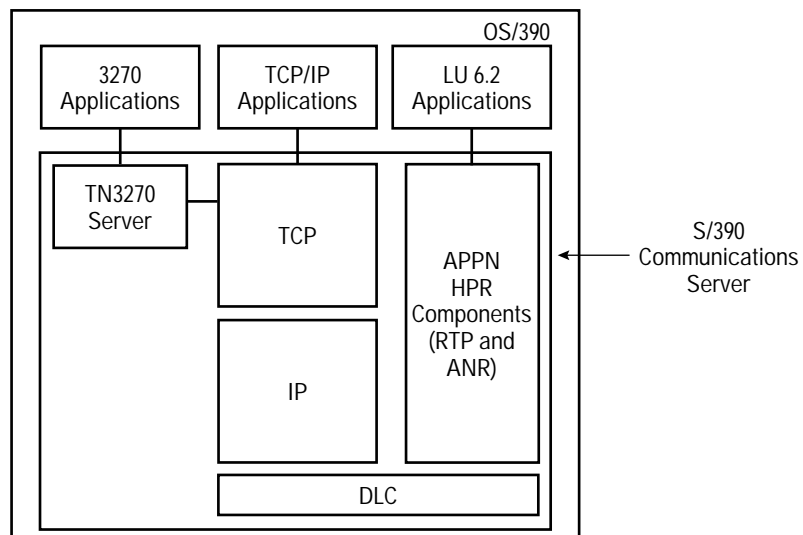
Figure 7: SNASw EE Network



As stated, the network is purely based on IP. More specifically, the SNA and APPN data is sent within User Datagram Protocol (UDP) packets, which is an IP-related standard data link control protocol. These packets are seamlessly transported across the network from end to end. The selection of UDP as the data link control was critical in providing prioritization of SNA traffic. With an EE solution, SNA traffic can be prioritized ahead of file transfers, e-mail, or other non-time-critical data, preventing SNA session loss.

The IBM S/390 Enterprise Server or IBM zSeries server is at the heart of an SNASw EE solution. The Communications Server software on these platforms (Communications Server for OS/390 and z/OS Communications Server, respectively) includes both TCP/IP and SNA/APPN networking logic. The newer releases of Communications Server also support EE, allowing HPR traffic to be sent over IP. All applications that are written to use APPN can automatically and without change take immediate advantage of EE. Older applications that use traditional 3270 datastreams can also be supported, by leveraging the TN3270 server included with Communications Server. Figure 8 depicts the architecture of today's Communications Server and illustrates how it supports native SNA, native APPN, and native TCP/IP applications.

Figure 8: Communications Server with EE Support





At the other end of the network, an EE-capable router or server accepts the UDP packets and interfaces them to an internal HPR stack for processing. HPR, at both the host and the client ends of the network, provides all of the error detection and correction logic necessary. In the case of traditional SNA devices such as controllers, terminals running emulators, and specialized branch equipment, the EE-capable router provides a function known as Dependent LU Requester (DLUR), which provides the capability for this traditional SNA traffic to participate in the EE design.

SNASw EE provides many benefits:

- Native IP from end to end, allowing SNA applications to take advantage of advances in IP technology
- Highly scalable solution that avoids early APPN broadcast issues
- Native SNA routing, allowing the removal of FEPs from the network
- Support for all types of traditional SNA devices and application traffic
- Capability of HPR to route around failures in the network, minimizing the risk of loss of SNA sessions
- Prioritization, allowing SNA traffic to take priority in the network over other data that is less time-sensitive
- Flexibility with support for a wide variety of devices, protocols, WAN speeds, and LAN media
- Enhanced flow control mechanism that allows EE servers to monitor and respond to congestion in the network and avoid packet loss
- Enhanced addressing with support for more than 64,000 network addressable units in a single VTAM domain (a limit in subarea SNA)
- Based on industry standards developed by the AIW and documented in IETF RFCs
- Superior usability, serviceability, scalability, and management when compared to first-generation APPN

SNASw EE is a solution created after almost a decade of experience implementing evolving APPN technologies in a wide variety of networks. SNASw EE and its sibling, BX, together enable organizations to build scalable, manageable, and resilient networks that preserve the investment in SNA/APPN devices and applications while bringing the advantages of IP throughout the network.

Implementation Scenarios

How an organization achieves the goal of an end-to-end logically integrated network depends on the point at which it is starting. In this section, several options for implementing this design are explored based on different existing situations. In all cases, the following environment is assumed:

- Multiple S/390 Enterprise Servers located in one or more data centers with recent versions of the operating system installed (OS/390 V2R6 or above); SNA, APPN, and TCP/IP applications currently running
- Multiple FEPs installed; SNA sessions are routed across several enterprise servers
- Large network supporting 25,000 sessions
- Multiple branch offices with requirement to support both SNA and TCP/IP applications and devices


These assumptions are made for the purposes of illustration only and to highlight the benefits of the various solutions. An organization whose current infrastructure differs from the illustrated environment may still benefit considerably from SNASw EE and related solutions; contact your Cisco sales representative or authorized partner to explore further.



SNASw EE

In the first scenario, the organization has not yet logically integrated SNA and TCP/IP within the network (that is, either physically separate or physically integrated stages). Figure 9 represents a typical customer environment at the physically integrated stage.

Figure 9: Typical Physically Integrated Environment



The branch offices contain SNA devices, TCP/IP devices, and a combination bridge and router that routes TCP/IP traffic and bridges SNA traffic. At the data center, the SNA traffic is bridged to a data center switch that is Token Ring-attached to two FEPs. The duplicate TIC addressing feature has been implemented, so the FEPs are able to act as hot standbys for one another. The TCP/IP traffic traverses the Ethernet-based campus backbone; TCP/IP traffic that is destined for the S/390 Enterprise Server arrives there using an older 3172 Interconnect Controller.

There are a number of problems with this current environment, including:

- The lack of SNA/IP integration results in a complex and difficult-to-manage environment with excess networking hardware, software, and management tools
- The eventual demise of FEPs makes planning for their replacement a high priority
- The Token Ring switching infrastructure needs to be migrated to Ethernet to provide a single infrastructure and to set the stage for future enhancements such as gigabit Ethernet (Note: Cisco has announced the End of Sale of its Token Ring switches)
- The 3172 Interconnect Controller, which is no longer sold, needs to be replaced with an updated device


This environment is the perfect target environment that would gain enormous benefits by adopting SNASw EE throughout the network. Starting at the branch locations, the existing bridge/routers would be updated with EE support. The DLUR capability of EE can support existing SNA devices, such as branch controllers and terminal emulators.

An alternative for branches that have servers based on CM/2 or PCOM is to place the EE functionality within these servers. However, there are several drawbacks to this alternative: duplication of networking function because the routers are still needed, multiple points of network management (router plus all servers), server CPU overhead, and increased cost and effort of upgrading each branch. For these reasons, Cisco recommends that the EE function and all networking function are concentrated on the branch router.



At the data center, the environment is simplified enormously by eliminating the FEPs and the Interconnect Controller and replacing them with either a Cisco router equipped with a channel connection or an IBM OSA-Express. Both platforms support high-speed TCP/IP connectivity to the S/390 Enterprise Server, and both can be implemented to provide redundant paths to the host. Figure 10 depicts the resulting solution.

Figure 10: SNASw EE Solution



The benefits of the new environment are:

- IP end to end
- Fewer networking devices, reducing both cost and complexity
- Preservation of SNA session routing after the replacement of the FEPs
- Token Ring migration to Ethernet can take place at the organization's own pace and can eliminate the dependence on duplicate TIC addresses for redundancy in the data center
- Optimal SNA session priority from the host all the way to the end devices
- Support for all existing devices and applications
- High availability solution; routing around failed components is automatic and seamless


This new network is now set to take advantage of all enhancements to IP technologies. It provides the foundation for the future, while continuing to support the vast investment already made in SNA and APPN devices and applications.

SNASw EE with DLSw+

In the second scenario, the network has already achieved logical integration (SNA over IP transport) on the backbone by deploying DLSw+ in the branches and in the data center routers. Figure 11 represents a typical Cisco customer in this scenario.



Figure 11: Typical Cisco DLSw+ Environment



As in the previous scenario, the branch offices have both SNA and TCP/IP devices. However, rather than simply bridging the SNA traffic to the data center, in this scenario the branch routers are equipped with DLSw+ and encapsulate the SNA data in IP. In the data center, one or more DLSw+ routers terminate the DLSw+ connection and bridge the SNA traffic to the existing FEPs. As in the previous scenario, the duplicate TIC addressing feature has been implemented, so the FEPs are able to act as hot standbys for one another. The TCP/IP traffic traverses the Ethernet-based campus backbone; TCP/IP traffic that is destined for the S/390 Enterprise Server arrives there using an older 3172 Interconnect Controller.


Some of the problems of this current environment include the following:

- SNA/IP in the backbone provides enhanced manageability; however:
 - The DLSw+ data center router is a single point of failure for SNA sessions
 - The scope of IP is limited to the backbone rather than all the way to the S/390 Enterprise Server
 - SNA session prioritization is mapped within the backbone rather than end to end
- The eventual demise of FEPs makes planning for their replacement a high priority
- The Token Ring switching infrastructure needs to be migrated to Ethernet to provide a single infrastructure and to set the stage for future enhancements such as gigabit Ethernet (Note: Cisco has announced the End of Sale of its Token Ring switches)
- The 3172 Interconnect Controller, which is no longer sold, needs to be replaced with an updated device

In this scenario, the migration to a more IP-centric environment can begin by updating the data center. The existing data center DLSw+ routers can be easily upgraded to support SNASw EE with DLUR support. With this addition, the DLSw+ connections are terminated by EE, and the traffic upstream of these routers is IP instead of SNA, allowing the elimination of the FEPs and the Interconnect Controller and, as in the previous scenario, replacement of them by either a Cisco router equipped with a channel connection or an IBM OSA-Express. Figure 12 depicts the resulting solution.



Figure 12: SNASw EE with DLSw+ Solution



The benefits of the new environment are:

- IP end to end
- Fewer networking devices, reducing both cost and complexity
- Preservation of SNA session routing after the replacement of the FEFPs
- Token Ring migration to Ethernet can take place at the organization's own pace and can eliminate the dependence on duplicate TIC addresses for redundancy in the data center
- SNA session priority maintained throughout the network
- Support for all existing devices and applications
- High availability solution; routing around failed components is automatic and seamless

This step of migration is a relatively easy change that minimizes the disruption to the network. After all, if DLSw+ has already been implemented throughout the network and has been tuned for the particulars of the environment, the least disruptive path is to leave that portion in place and focus on updating the data center. However, an organization in this scenario may decide eventually to migrate to the pure SNASw EE environment described in the previous scenario. The reasons for considering this long-term migration are:

- Reduced complexity
- Better end-to-end SNA session prioritization
- Better support for bridged Ethernet environments
- Elimination of the DLSw+ endpoint in the data center, which is a potential point of failure
- Better end-to-end session redundancy by leveraging host-based session routing capabilities

The combined SNASw EE with DLSw+ solution is a proven and solid solution, implemented by many Cisco customers. It can stay in place for many years, allowing organizations to migrate to a pure SNASw EE environment at their own pace.



Getting Started

The withdrawal of the IBM 3745/3746 Communication Controller has provided the impetus for all organizations that continue to support SNA traffic within the network to put in place a plan to migrate to an IP-based solution. Fortunately, Cisco, IBM, and their respective partners can offer a great deal of assistance in formulating and implementing the migration plan.

The first step is to carefully inventory the existing environment. Particular attention should be paid to the precise number, location, and configuration of FEPs. To assist in this effort, IBM has produced the *IBM Communication Controller Migration Guide*, a redbook that is available from IBM directly or on its Web site. This comprehensive guide provides tools and information to allow organizations to consolidate and migrate from their existing FEPs.

The next step is to define the near-term and long-term goal infrastructure. Because native SNA routing within the network is eventually going away, it is critical that IT organizations determine a stepwise plan to move to IP. Identification of a timeline with milestones is critical, as is prioritization of projects. The *Cisco SNA Internetworking Design and Implementation Guide* is a very good reference to assist in this planning.

One important step in the migration, for those organizations with SNA/APPN-only S/390 servers, is to identify a plan to enhance the S/390 server complex to support IP. It is not necessarily a requirement that all S/390 servers directly support IP, but those that are interfacing to the network must eventually be upgraded to include IP support. For some organizations, this is a big step that must be carefully planned and implemented over time.

If it will be some time before the S/390 Enterprise Servers are equipped to support IP, the organization may consider beginning the logical integration of SNA and IP on the backbone using DLSw+ if that step has not already been taken. This provides an interim step to the eventual goal of IP from end to end. SNASw EE provides the ultimate solution to meet this eventual goal.

An IT organization does not need to embark on this path alone. Cisco, IBM, and their partners can offer a great deal of assistance. Together, they have helped thousands of companies worldwide along the migration path. That experience and expertise has resulted in numerous case studies that demonstrate actual customer experiences and a vast amount of data showing before and after results. To get started, many documents are available on the Cisco and IBM web sites for further information.

Cisco Documents

- Cisco SNA Switching Services Web page:
<http://www.cisco.com/warp/customer/cc/pd/ibsw/snasw/>
- *SNA Switching Services Design and Implementation Guide*:
http://www.cisco.com/warp/customer/cc/pd/ibsw/snasw/tech/snst_rg.htm
- *SNA Internetworking Design and Implementation Guide*:
http://www.cisco.com/warp/public/cc/pd/ifaa/ifpz/chifpz/tech/dcdgr_rg.pdf
- “Extending the Enterprise: SNA Application Transport over IP” white paper:
http://www.cisco.com/warp/customer/cc/so/neso/ibso/ibm/s390/eesna_wp.htm

IBM Documents

- IBM Communication Controller Web page:
http://www.ibm.com/servers/eserver/zseries/networking/374x_resources.html
- IBM 3745/3746 Communication Controller announcement letter:
<http://www.networking.ibm.com/announce/022602.html>
- *IBM Communication Controller Migration Guide*:
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246298.html>
- “Enterprise Extender” white paper:
<http://www.ibm.com/software/network/library/whitepapers/eextender.html>
- “Enterprise Extender: A Key to SNA/IP Integration” paper:
http://www.ibm.com/servers/eserver/zseries/library/techpapers/enterprise_extender.html



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2002, Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, and the Cisco Systems logo are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0206R) SPS 08/02