


IBM COMMUNICATIONS CONTROLLER REPLACEMENT DESIGN GUIDE

Much of today’s business-critical applications and data remain on IBM mainframes. While many SNA-based applications continue to run on the OS/390 or z/OS operating systems, most new mainframe-based applications are based on IP. Almost 90 percent of network access to both existing and new applications has moved to TCP/IP either natively or utilizing SNA/IP solutions (see Figure 1).

Figure 1

WAN Connectivity to Mainframe Servers




Source: IDC, Mainframe Networking Opportunity Analysis, October 2002

Given these changes, higher bandwidth and robust IP networking to the mainframe server are required. At the same time, IBM has discontinued sales of its 3745 Communications Controller and 3746 Nways Multiprotocol Controller, also known as front-end processors (FEPs). Customers are looking for networking solutions that continue to support their existing mainframe-based business requirements while enabling their mainframe for its new role as an IP server.

IBM’s SNA has been the mainframe’s network protocol of choice for more than 30 years. As recently as 1998, International Data Corporation (IDC) estimated that more than 70 percent of all mission-critical applications in the world run on IBM mainframes in large data centers using SNA. However, with the shift to IP as seen in Figure 2, the mainframe environment is evolving.

Figure 2
Network Protocol Trends



Source: IDC, Mainframe Networking Opportunity Analysis, October 2002

IBM, from a server perspective, has been working to evolve the zSeries to support a robust IP and Web environment. It has added Linux support to the mainframes and optimized the TCP/IP stack in z/OS. It also added high-speed LAN access to the mainframe to accommodate the increased bandwidth requirements of the Web-based and IP applications. This shift to IP also has led IBM to discontinue sales of the FEP.

This paper begins with an overview of the transformation of the information system from traditional mainframe protocols (SNA, Binary Synchronous Communications [Bisync], and so on) to IP. After a quick review of the FEP features, the paper details the functions Cisco provides when replacing the FEP with connectivity from either the channel-attached router or the Gigabit Ethernet-attached switch, as well as support at the edges of the network for legacy devices.

This paper also compares the functions provided by a Cisco® channel-attached router and Gigabit Ethernet-attached Cisco Catalyst® switch to the features and functions provided by the IBM 3745/3746 FEP in the areas of:

- WAN connectivity
 - Synchronous Data Link Control (SDLC)
 - X.25 and Qualified Logical Link Control (QLLC)
 - Frame Relay, boundary network node (BNN), and boundary access node (BAN)
 - Bisync and async
 - Airline protocols
- Connectivity to the LAN and data center backbone
 - Token Ring
 - Ethernet
 - Gigabit Ethernet

- SNA Support
 - SNA routing (subarea and Advanced Peer-to-Peer Networking [APPN]) and bridging (Data Link Switching Plus [DLSw+])
 - Boundary function for subarea-based endpoints
 - Dependent Logical Unit Requester (DLUR)/Dependent LU Server (DLUS) support
 - Interconnect between SNA networks (SNI)
- Mainframe physical connectivity

Please note that Cisco solutions for FEP migration do not cover all FEP functions and features. Those that are not covered are listed at the end of this paper.

SNA TO IP TRANSFORMATION OF THE DATA CENTER

With the migration toward Web-based applications underway, it is time to rearchitect the network to provide optimum support. With the application rollout often requiring changes at the server, the client, and the network, it is a multiphase process. Phases requiring network changes must also maintain the network's stability, availability, and security for the applications currently supporting the business. SNA/IP integration may be part of a larger process to upgrade the network for all communication needs (data, voice, and video). This network upgrade may be implemented in stages:

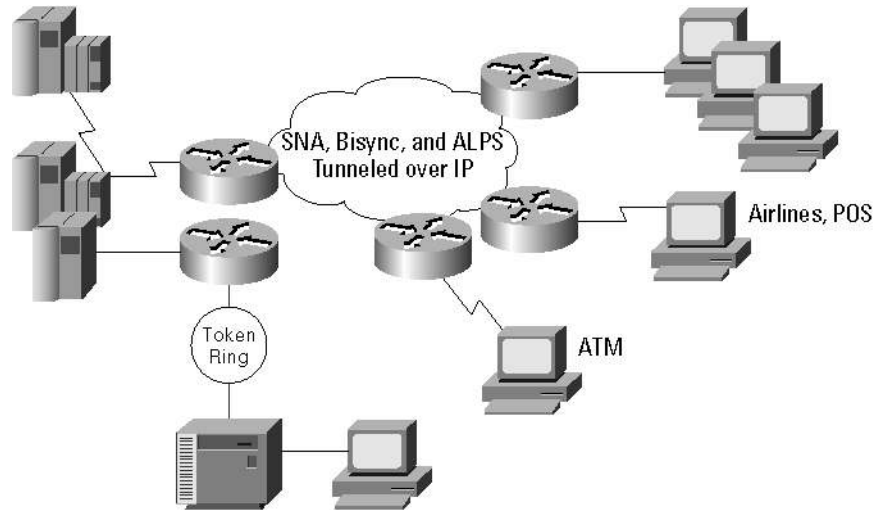
- Robust IP network infrastructure and WAN consolidation
- Data center SNA infrastructure upgrade including multiprocessor, multisite, and storage
- Branch transformation of endpoints including wireless and IP telephony

Because the replacement of the FEP is forcing change in the network, many customers are taking advantage of this to rearchitect their network to handle their future needs for network bandwidth, reliability, availability, and operations as well as integrate new technologies like voice, video, and wireless. This evolution, rather than revolutionary change, enables support of current applications and connections to customers, employees, suppliers, and partners, while providing the infrastructure required by new and emerging e-business applications.

Robust IP Network and WAN Consolidation

Many organizations accomplished the first stage, WAN consolidation, in the mid- to late 1990s, encapsulating SNA and other protocols over IP and Frame Relay networks (see Figure 3). Cisco provides many tunneling technologies for this stage including DLSw+, serial tunnel (STUN), block serial tunnel (BSTUN), Airline Product Set (ALPS), and Frame Relay access support (FRAS).


Figure 3
SNA to IP WAN Consolidation



Data Center Infrastructure Update and Support for Web Applications

With the additional roles that the mainframe is taking on, it must maintain and, in some cases, enhance the level of business resiliency that customers came to expect from the SNA-based environment. With the advent of optical interconnect and Storage Area Networking in the data center, backup and recovery policies, procedures, and implementations can become consistent across all servers. Mainframes require high-speed connectivity to participate fully as an IP server. IBM Open Systems Adapter (OSA)-Express network interface cards (NICs), when combined with the Cisco Catalyst 6500 Series Gigabit Ethernet switch, support high-speed connectivity for TCP/IP applications on the mainframe. Cisco routers with the Enterprise System Connection (ESCON) Channel Interface Processor (CIP) or Channel Port Adapter (CPA) can also support TCP/IP (as well as SNA) over a multitude of LAN and WAN choices in this environment (see Figure 4). More information on Data Center Networking can be found at <http://www.cisco.com/go/datacenter>.

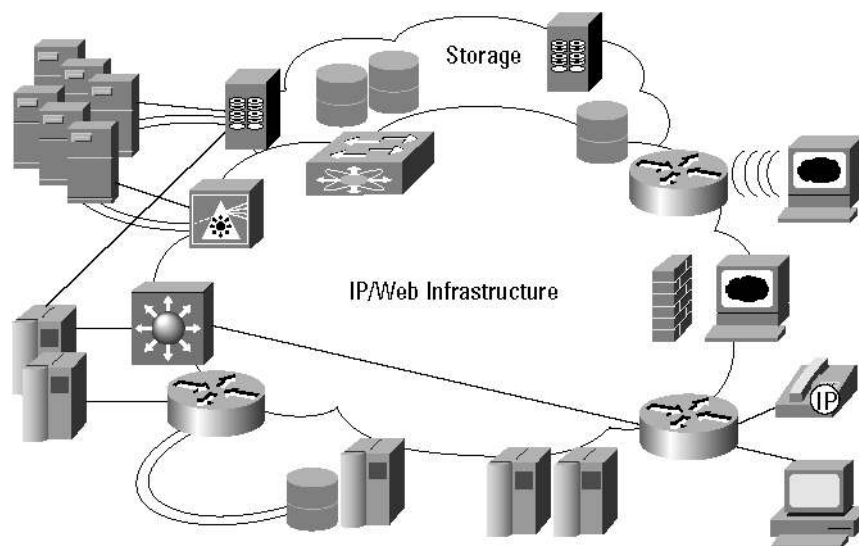
Figure 4
Updated Data Center Infrastructure



Branch Transformation

The final step of this evolution is to move toward branches that have only devices based on IP. Of course, this cannot always happen immediately so the utilization of some SNA/IP technology may be required as well. This “transformed” branch can then take advantage of the mobility that wireless connections provide and the dynamic voice support of IP Phones (see Figure 5).

Figure 5
Branch Transformation



More information on branch transformation can be found at The Business Ready Branch Networking Solution at http://www.cisco.com/en/US/netsol/ns477/networking_solutions_packages_list.html and the Cisco/IBM Branch Transformation home page at http://www.cisco.com/en/US/partners/pr67/pr30/pr216/partners_strategic_solution_concept_home.html.

Today, with networks changing rapidly to meet business demands, the typical enterprise network must support highly available IP solutions over LANs and WANs, as well as other services such as telephony. High-performance connections such as LANs, high-speed serial lines, and Frame Relay have replaced low-speed serial lines. The FEP's hardware, designed for low-speed serial concentration, has not kept up with the high bandwidth and open systems requirements of today's enterprise networks. As a result, additional networking gear is required to augment or replace FEPs in these environments.

The FEP chassis sales have declined to the point where IBM has now removed them from their sales offerings. Many IT professionals are in the process of determining the best replacement solution to connect their mainframes and servers, balancing the requirements of their current applications and end users, while positioning their systems for the business opportunities becoming available with e-commerce.

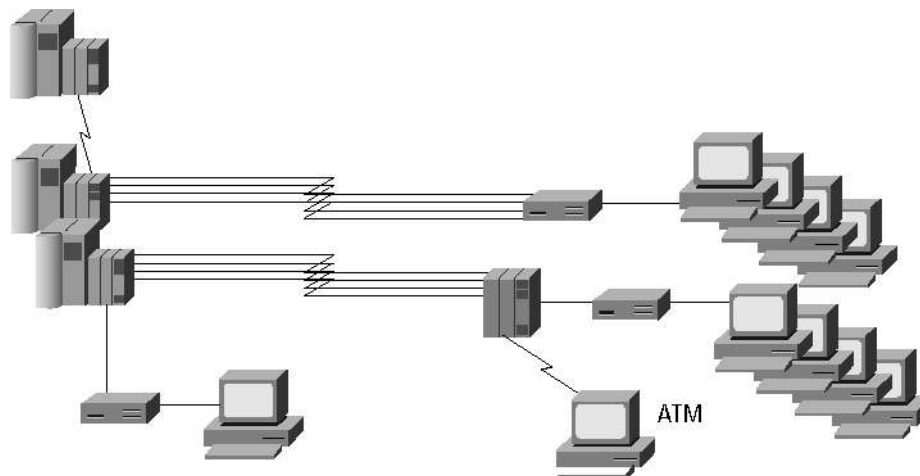
If an organization is considering replacing some or all of its FEPs, the first step is to determine the functions, features, and environment that the FEP is supporting today. With those requirements, features can then be mapped as the network moves to high-speed, multifunction solutions such as CIP- and CPA-attached routers or OSA-Express-attached switches. In some cases it will be possible to move specific functions into a mainframe, by running the FEP software in a hardware emulation mode on top of Linux. This recent solution is being made available in stages to accommodate some of the unique functions of the FEP.

REVIEW OF IBM 3745/3746 FEP

Historically, IBM's primary solution for SNA mainframe access has been the FEP running Network Control Program (NCP) software. From the late 1970s, the IBM FEP attached directly to the mainframe or at remote sites provided SNA subarea functionality and legacy protocol support (see Figure 6). Only the largest networks use the full range of functionality provided by the FEP; most small networks use only a subset. FEPs serve key functions in today's networks in the areas of WAN concentration, LAN connectivity, SNA routing and support, legacy application support, high availability, and physical connectivity.

Figure 6

FEP Placement in the SNA Network



WAN Concentration

FEPs provide the following WAN concentration functions:

- *Leased Line SDLC*—FEPs can concentrate large numbers of low-speed (9.6-kbps) serial lines. However, as networks migrate to higher-speed switched WAN backbones, the need for high-density, low-speed serial connectivity decreases.
- *Switched SDLC*—Some enterprises rely on switched SDLC to support transient SNA connections to small branch offices or to provide switched network backup.
- *X.25 support*—X.25 Interconnection allows the NCP to act as an X.25 packet switch. NCP Packet Switching Interface (NPSI) allows the NCP to connect to other resources over X.25 networks. X.25 Interconnection supports both SNA and non-SNA devices. For non-SNA (async and Bisync) devices, it supports conversion to SNA.
- *Frame Relay*—NCP supports Frame Relay as a Frame Relay data terminal equipment (DTE) device or a Frame Relay switch.
- *Legacy protocols*—The FEP supports program products, such as Non-SNA Interconnection (NSI) for Bisync conversion, Airline Line Control Interconnection (ALCI) for airline line control protocol transport, and Network Terminal Option (NTO) for asynchronous/bisynchronous conversion. Bisync continues to be used by some devices and applications for which the FEP provides a conversion function, so that they can be treated as SNA devices in the host.

LAN Connectivity

The FEP supports the following LAN connectivity options:

- *Token Ring*—Token Ring was added as the first LAN connection to the FEP in the 1980s. Many customers used Token Ring as their data center backbone. Token Ring's ability to logically support duplicate Media Access Control (MAC) addresses in a source-route bridged network added high availability for SNA in the data center.
- *Ethernet*—Ethernet was added to the FEP as an additional connection in the 1990s. Given its inability to support duplicate logical MAC addressing in a transparent bridged environment, Ethernet did not catch on as a FEP connection for SNA.

SNA Routing and Support

The FEP provides the following types of support for SNA routing:

- *SNA session routing*—SNA session routing is required in environments with multiple data centers or Advanced Communications Function (ACF)/Virtual Telecommunications Method (VTAM) application hosts and a high volume of cross-subarea SNA traffic. This is frequently seen in what is called the Communication Management Configuration (CMC) design model, where VTAM is assigned to "own" the network resources and other systems are tasked with running business applications.
- The FEP also worked in concert with VTAM to provide Composite Network Node support for APPN and subarea networks.
- *SNA Class of Service (COS)*—SNA COS enables prioritization of SNA traffic between the FEPs and the mainframes and is important in environments with SNA backbones. SNA COS is less important in environments that have consolidated the FEPs in the data center. In this case, either there is no FEP-to-FEP traffic, or the FEPs are connected at the data center over high-speed LANs that do not have bandwidth contention problems. However, some networks take advantage of Link Services Prioritization (LSPRI), which provides transmission priority based on COS for outbound traffic (for example, FEP to cluster controller).
- *SNA BNN function*—FEPs provide an SNA BNN function, which includes polling, converting from local addresses to SNA addresses, and converting exchange identification (XID). In the absence of remote FEPs, local FEPs can perform these functions. In the absence of any FEPs, VTAM can perform most of these functions.

- *SNI*—Most enterprises use FEPs for SNI to allow independent SNA networks to communicate. There are other alternatives, such as electronic data exchange over the Internet or intranet, the APPN Extended Border Node (EBN) function over IP (Enterprise Extender [EE] or High Performance Routing [HPR]/IP), or EBN using traditional Layer 2 protocols via a DLSw+ network. It is important to note that for most of these alternatives, both ends of the connection must change. Single-sided adjacent SNI connections are also being used when the partner organization must keep its FEP or will not move to an IP-based solution.

Legacy Application Support

FEPs offer specialized program products that support custom or older applications. Network Routing Facility (NRF) provides routing inside the NCP without VTAM participation. An emulation program allows the IBM 3745 to connect to Basic Telecommunications Access Method (BTAM) in an IBM mainframe.

High Availability Options

The FEP provides the following high availability options:

- *System services control points (SSCP) takeover*—With this facility, if an owning VTAM goes down, another VTAM can assume ownership of those resources without disrupting any existing application sessions. The NCP plays a role in allowing this takeover.
- *Extended recovery facility (XRF)*—XRF is a program that allows one VTAM application to take over for another. The XRF code in the NCP plays a key role in supporting this capability. IP networks designed for high availability help to alleviate this issue.

Physical Connectivity

FEPs provide the following physical connectivity options:

- *ESCON*—The FEP connects to the mainframe via ESCON using the Channel Data Link Control (CDLC) protocol.
- *Bus and tag*—Some installations may still be supporting parallel port (bus and tag) connections.

USING CISCO SOLUTIONS AS A FEP ALTERNATIVE

To what extent can a Cisco router replace a FEP? The first, and most basic, function of a FEP is to offload the work of handling serial connections. Cisco routers do an excellent job of terminating SDLC or Frame Relay serial connections of any speed, and then passing that traffic across a channel to the mainframe. Cisco routers also handle LAN connections very well, usually Token Ring, using Logical Link Control, type 2 (LLC2) as the data link control layer protocol and terminating LLC2 with either DLSw+ or directly on a channel adapter. The second most important function of FEPs, SNA routing, is handled by Cisco SNA Switching Services (SNASw). Beyond these basic FEP functions, some of the 200 plus other functions can be replaced by some of the 400 plus functions of the Cisco IOS® Software.

Cisco has two primary designs for FEP migration: channel-attached router (Cisco 7500 Series with CIP or Cisco 7200 Series with ECPA4) or Gigabit Ethernet-attached switches (Cisco 6500 Series with IBM OSA-Express). Both types of mainframe connectivity work in conjunction with Cisco IOS-based features such as SNASw and DLSw+ to provide similar features and functions of the IBM 3745/3746 FEP.

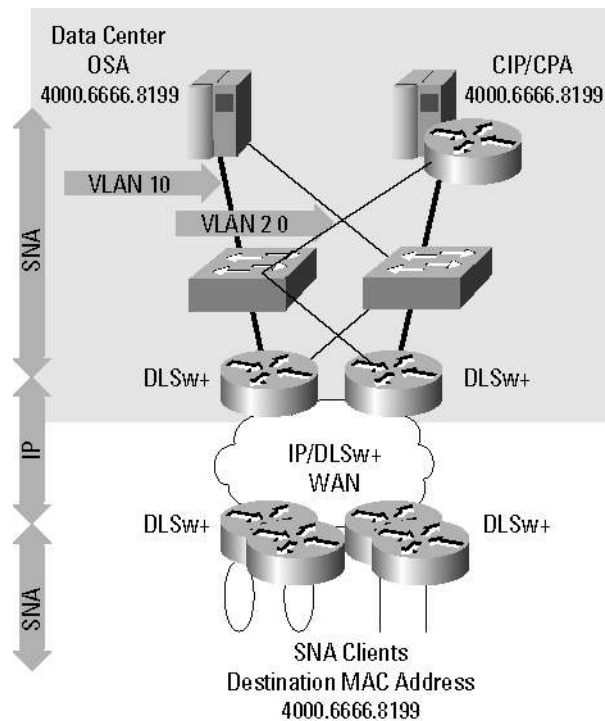
LAN Connectivity

The channel-attached router supports all standard LANs, including Token Ring and Ethernet like the FEP does. It also supports Fiber Distributed Data Interface (FDDI), Asynchronous Transfer Mode LAN Emulation (ATM LANE), Fast Ethernet, and Gigabit Ethernet.

Given that Token Ring network devices are becoming more scarce and that network architects are no longer limited by the set of LAN connections provided by the FEP, organizations may decide to enhance their data center backbones with a higher-speed, highly available infrastructure based on Gigabit Ethernet. When moving from a Token Ring-based network that supports duplicate Token Ring interface coupler

(TIC) MAC addresses as an SNA high-availability feature, the new solution must also support high-availability features for both IP and SNA. For example, SNASw has Hot Standby Router Protocol (HSRP) support for LAN connectivity, enabling a backup SNASw router to take over the network MAC address of the primary SNASw router. For the bridged SNA network, built with DLSw+ and the CIP, the solution is to configure duplicate TICs or run the CIP or CPA in hot standby. Duplicate MAC addresses can be used with Ethernet, as long as there is no transparent bridge connecting them. This makes it possible to design a DLSw+-based network that includes redundant DLSw+ headend routers, which access separate virtual LANs (VLANs) that each include a duplicate Ethernet MAC address (see Figure 7).

Figure 7
DLSw+ with Duplicate Ethernet MAC Addresses in Head-End Routers



As Figure 7 shows, this can be accomplished using separate VLANs, taking care not to allow any Layer 2 connectivity between these VLANs. This design still retains the benefits of DLSw+:

- Up to four active peer paths cached
- Customization using peer costing
- Load-balancing capability, with multiple circuit balancing options

Although Layer 2 bridging designs can be created to roughly duplicate Token Ring capabilities, the preferred solution is to move to EE for SNA transport, providing an even higher level of availability. For more information on Token Ring-to-Ethernet migration, see

http://www.cisco.com/en/US/netsol/ns340/ns394/ns74/ns149/networking_solutions_package.html.

WAN Connectivity

Communication lines operate in half- or full-duplex mode, using SDLC, Point-to-Point Protocol (PPP), Frame Relay, or X.25 line protocol. All of the physical interfaces supported by the IBM 3745 are supported by Cisco routers. Cisco supports SDLC, QLLC, Bisync, async, and the airline protocols over serial adapters. These serial adapters also support Frame Relay, X.25, High-Level Data Link Control (HDLC), and PPP.

Serial Concentration

Serial concentration can be supported via the low-speed serial interfaces across the line of Cisco routers. At the high end, the Cisco 7200 Series router (which also supports the CPA) supports 48 serial lines and one LAN. For aggregating WAN connections in a separate router, the Cisco 3600 Series router supports up to 48 serial lines, or 36 serial lines plus one LAN. The least expensive solution is buy Cisco 2600 routers, then “just rack ‘em and stack ‘em.”

When converting serial SDLC to LAN LLC2, to be bridged via DLSw+ or routed via SNASw, the SDLC protocol is terminated at the SDLC-attached router. Transaction response times may actually improve with the right IP prioritization scheme in the private intranet. This produces very nearly the same response time characteristics as a directly attached serial line. Usually there are faster lines in the IP network, and serialization delay on low-speed lines accounts for the majority of delay. Therefore, most customers see faster response times.

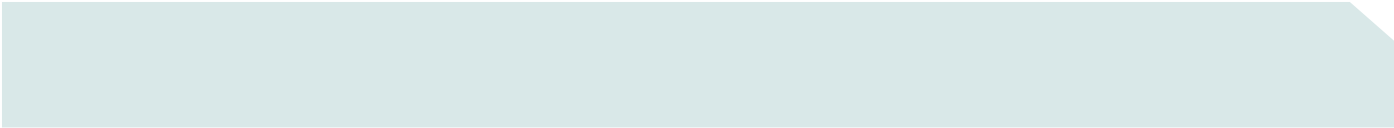
Multidrop SDLC connections are supported by DLSw+. It will increasingly make more sense, as carriers deploy virtual private networks (VPNs) and Multiprotocol Label Switching (MPLS) networks, to terminate the SDLC circuits at the remote location and transport them through an IP network.

X.25 and start-stop or async lines can be converted to Telnet sessions, and these sessions can be connected to a custom TCP/IP application on the mainframe. There are two different cases for Bisync. Interactive 3270 connections should generally be replaced by TN3270 clients and servers. In the case of Bisync automated teller machines, Cisco can convert the Bisync data to formatted TCP/IP data, which can be easily terminated by any TCP/IP application on the mainframe. Batch 3780-type traffic should be terminated on one or more separate machines equipped with one of the many Bisync emulation packages. The data can then be easily transferred to the mainframe using either a TCP/IP application such as File Transfer Protocol (FTP), or an SNA emulator using the LAN for connectivity.

Switched SDLC

The Cisco IOS Software transports multiprotocol traffic, including SNA, over switched services. However, it does not directly support dial-out to switched SDLC devices. If the SDLC dial-up connections are inbound to the Cisco router via a sync modem, the router raises data terminal ready (DTR), and when the modem accepts the call, it raises data set ready (DSR). The line then comes up and starts communication. DLSw+ would then be used to transport the SDLC.

Note: Dial-in requires coding `sdlc role prim-xid-poll` on the appropriate serial interface.



Support for “dial out” SDLC is quite limited. If there is a single number that needs to be dialed, it can usually be stored in the modem, and when the associated VTAM resource is activated, the stored number is dialed. The method where the phone numbers are stored in VTAM and sent to the modem via V.25 is not supported by Cisco. One alternative is to have devices that are SDLC-attached to a Cisco router defined as nonswitched and then have a switched (dialup) IP network between the endpoints. Using Dial on Demand for IP or an Internet connection at the branch, SDLC could be supported if there is a router at the remote site. Making a change to pure IP for these solutions is also a choice. With the IP solution, all updates could be done in parallel.

Frame Relay

Frame Relay, including BAN and BNN, is supported in the Cisco router. Where direct attachment to IBM devices is required, the RFC 1490 protocols can be used to match the encapsulations and provide a low-overhead bridging function. However, customers typically use IP over the Frame Relay network and then use DLSw+ with TCP encapsulation for maximum availability and flexibility.

QLLC


Cisco can terminate QLLC connections and recommends using DLSw+ to do this. QLLC can then be converted to LLC2 or, optimally, be handed off to SNASw to be routed to the mainframe using EE.

Legacy Protocols

Although the Cisco IOS Software can duplicate some special protocols supported by the FEP, such as async and Bisync tunneling, comparable protocol support (that is, conversion to SNA) is not provided. For Bisync, Cisco routers can encapsulate various Bisync protocols (3780, 3270, and generic) into IP and transmit over a WAN using BSTUN. This requires Bisync at both sides.

When removing a FEP supporting Bisync lines that connect many automatic teller machines, traditional BSTUN is no longer useful, because the serial interface is gone. To avoid immediately replacing or upgrading these existing devices, Bisync over IP (BIP) allows a single-sided 3270 Bisync connection where the serial Bisync is converted to IP, which is then delivered to the mainframe over CIP, CPA, or Gigabit Ethernet connections (see Figure 8). The limitation for BIP is that the Bisync 3270 control unit can contain only a single 3270 device.

Figure 8
Bisync to IP Conversion



Most 3780 Bisync is used for file transfers and should be replaced by an IP-based application. Several large corporations have mandated the use of Applicability Statement 2 (AS2) as described by the Internet Engineering Task Force (IETF) Electronic Data Interchange-Internet Integration (ediint) group (see <http://www3.ietf.org/proceedings/05nov/ediint.html>). Many other organizations are using the readily available FTP, which can be easily integrated into automated processes. If the network must continue to support batch Bisync lines, the best approach is to select from the numerous vendors who provide a standalone Bisync host function on PC or UNIX platforms.

X.25 Support

The Cisco IOS Software can be configured as an X.25 packet switch and supports transport of SNA over an X.25 backbone. However, Cisco has no comparable function to provide X.25 originating traffic conversion to SNA. Host-based software products available from Comm-PRO and Computer Associates provide much of the same protocol conversion capability of the NPSI product for FEPs, by working with Cisco X.25 over TCP/IP (XOT).


SNA Support

Cisco offers a variety of products and solutions for supporting SNA within an IP-based infrastructure.

SNA Session Routing

Cisco SNASw, a feature of the Cisco IOS Software, supports native APPN routing for both APPN and subarea SNA clients. It is an APPN node product that implements the IBM EE (RFC 2353) to support integration of APPN into the IP network and Branch Extender (BX) architectures to improve scalability (see Figure 9). SNASw is not a pure FEP replacement, but it will handle independent LUs as well as dependent LUs (via DLUR). SNASw provides routing of 3270 traffic and LU 6.2 over either an APPN or IP network. More dynamic and less labor-intensive to maintain than a subarea network, SNASw extends the number of SNA network addressable units (NAUs) beyond the 64,000 limit per subarea SNA domain.

Figure 9
SNASw



The HPR session path can be either SNA or IP (with EE turned on). The ESCON-attached router supports a native SNA path to the mainframe. When EE is used, the channel-attached Cisco router or Gigabit Ethernet-attached Catalyst switches can be used for higher bandwidth to the IP stack on the mainframe. The IP stack on the mainframe uses a virtual IP address (VIPA) as the IP addressable endpoint.

By using HPR/IP upstream from a router, the robustness and redundancy in the IP network can be exploited to transport SNA traffic to the host. It also offloads the LLC-to-IP conversion to the router. SNASw can accept SNA natively from LAN connections or work in conjunction with DLSw+.

In-depth design considerations such as placement, performance, naming/addressing, redundancy, and maintenance of an SNASw network is outlined in the SNA Switching Services Design and Implementation Guide at:

http://www.cisco.com/application/pdf/en/us/guest/tech/tk897/c1501/ccmigration_09186a008007ca58.pdf.

Operational and troubleshooting techniques can be found in the white paper, "Troubleshooting SNASw: A Guide to Problem Resolution," at:

http://www.cisco.com/en/US/tech/tk331/tk336/technologies_white_paper09186a008011ff1b.shtml.

TN3270 Server

The Cisco CIP runs TN3270 Server and connects via ESCON to the mainframe. TN3270 client traffic comes into the channel-attached router as IP and is converted by TN3270 Server to SNA/LLC to connect to the mainframe. The extension of the number of SNA NAUs is supported by the Cisco TN3270 Server when using the APPN DLUR capability. When the TN3270 Server is run on the channel-attached router, the only host MIPs used are in the SNA stack. If the TN3270 server function is actually running on the mainframe, then the MIP load is shared across the host IP and SNA stacks.

Bridging SNA with DLSw+

If APPN cannot be run in the data center, a channel-attached router running Cisco SNA (CSNA) channel function and DLSw+ can provide support for SNA traffic. All SNA routing would then be done on the mainframe. DLSw+ supports Token Ring- and Ethernet-attached LAN devices as well as SDLC and QLLC serial-attached devices.

SNA COS

SNASw also preserves SNA COS for both Advanced Program-to-Program Communications (APPC) and 3270 traffic. If SNASw is installed only in central site DLSw+ routers, it provides outbound prioritization based on COS, similar to LSPRI in the FEP. In a multiprotocol environment, one of the various Cisco queuing algorithms can be used to reserve bandwidth for SNA traffic over DLSw+ (DLSw+ also supports LU and service access point [SAP] prioritization). SNASw EE supports SNA COS-to-IP type of service (ToS) mapping (SNA transmission priority-to-IP precedence mapping) for both inbound and outbound traffic in a bidirectional fashion between an EE-enabled S/390 host and a SNASw EE router, allowing the service policy agent within CS/390 to set quality of service (QoS) policies, which will then be enforced throughout a Cisco IP network.

SNA BNN Functions

The Cisco IOS Software in a DLSw+ environment can reduce mainframe cycles by providing several boundary functions such as remote polling, group poll support, and downstream physical unity (DSPU) concentration. Using SNASw and DLUR, Cisco routers can provide all of the BNN functions provided by a FEP.

SNI

The SNI function is not supported in the Cisco IOS Software. The only entity that actually does SNI is NCP. Cisco provides the robustness and resiliency in its IP support needed to interconnect disparate SNA networks. If both the primary and partner company can agree, the most strategic move is to replace the SNI applications with extranet IP applications. Both companies can then take advantage of VPNs, security, routing availability, and the ongoing enhancements of the Internet and corporate intranet. As alternatives, EBN over EE or over LLC, Casual Connect, or single-side SNI can be implemented to continue the SNA applications exchange.

If both companies need to keep the SNA applications intact but can support APPN in both networks, they can still take advantage of the intelligent IP networking features by implementing EBN over EE. This function is performed by z/OS Communications Server, which allows SNA traffic to flow between two distinct SNA networks over an IP connection directly from the mainframe, providing dynamic paths and connectivity. This IP connection can run over the Gigabit Ethernet-attached switch or the channel-attached router. With APPN EBN and EE replacing SNI for both partners, EBNs on each side act as transit nodes. This is true unless Global Connection Network (also known as Global Virtual Routing Node [GVRN]) can be implemented. GVRN allows an EE connection (and hence session traffic) directly between two peers in different networks. The session establishment flows still pass through the EBNs, but the session does not. GVRN has a few restrictions that should be resolved in the next z/OS release. For example, GVRN does not work when it is adjacent to an interchange (APPN/subarea boundary) node, and it currently does not work with firewalls (that is, Network Address Translation [NAT]).

If the companies choose not to use EBN with EE, they can use EBN over LLC through the channel-attached router. They can also implement Casual Connect, which uses a low-entry networking (LEN) type connection. This requires hard-coded definitions of all the resources being used, but the FEP can be removed at either or both ends for this connection. CSNA in the Cisco router can be used to bridge the PU 2.1 traffic over the channel, and DLSw+ can be used to transport it over the WAN. Each resource that is to be referenced in the other network needs to be predefined. This only supports application-to-application traffic.

If the partner is not ready to support APPN or IP applications, one-sided SNI connections can be used. To do SNI requires at least one FEP, but a Cisco CIP can be used at the other side of the connection (see http://www.cisco.com/en/US/tech/tk331/tk336/technologies_configuration_example09186a0080093ec9.shtml). This would allow one company

to replace its FEP as long as there is still a FEP at the other end of the SNI connection. SNI requires only one gateway NCP and one gateway SSCP. The Cisco IOS Software allows connection to an SNI gateway, but it does not provide SNI gateway functionality. A Cisco router can be connected to this NCP either via SDLC or Token Ring. It can then transport the SNA across a WAN using DLSw+. At the DLSw+ endpoint it can then go directly to VTAM via a CIP router.

More recently IBM has created software that allows the NCP to run on a Linux operating system on the mainframe. Since this is the same NCP program that runs on FEPs, it also supports the GWNCP function for SNI. This solution; Communication Controller for Linux (CCL) and NCP, provides for physical attachment via Toke Ring Interface Couplers (TICs), which are them mapped to OSA adapters on the mainframe. For wide area network connections, and support of other media types and protocols, IBM has tested and recommends Cisco DLSw+.

SSCP Takeover

The SNASw DLUR feature of the Cisco IOS Software fully supports the SSCP takeover facility.

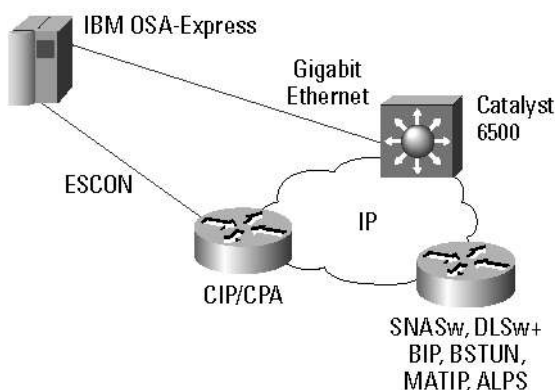
Mainframe Connectivity

Cisco solutions for physical connectivity to the mainframe (see Figure 10) include the following:

- Channel-attached router via the Cisco CIP or CPA with ESCON channel attachment with the following channel protocols:
 - CSNA supporting Link Services Architecture (LSA) for subarea and APPN SNA natively over the channel to z/OS
 - Common Link Access for Workstations (CLAW) or Packed CLAW supporting IP over the channel to either z/OS or Linux
 - Cisco MultiPath Channel Plus (CMPC+) (also called MPC+ by IBM) supporting multiple read/write paths for IP over the channel to z/OS
- Cisco Catalyst 6500 Series Gigabit Ethernet-attached switch with IBM OSA-Express providing high-speed LAN access (Gigabit Ethernet, Fast Ethernet, Ethernet, and ATM) for mainframe-based IP applications

Figure 10

Mainframe Connectivity Options



SELECTING THE MAINFRAME CONNECTION

When replacing the FEP, there are two ways of connecting to the mainframe: a Cisco 7500 or 7200 Series Router equipped with a CIP or CPA, and a Cisco Catalyst 6500 Series switch connected to an IBM OSA-Express.

Using a Cisco Channel-Attached Router Solution

A Cisco channel-attached router with a CIP or CPA offers benefits that are not available in an IBM 3745. These benefits include:

- *Multipurpose solution*—The CIP and CPA provide a state-of-the-art, high-performance solution for mainframe connectivity for access not only to SNA applications but to TCP/IP applications as well. As networks begin to offer intranet and Internet services, tying the mainframe into TCP/IP networks with features such as TN3270 Server enables organizations to leverage their mainframe investments. NCP's support of TCP/IP is limited.
- *Higher speed*—The CIP and CPA offer a tenfold improvement in performance over an IBM 3745 model 200 for SNA, and an even larger improvement for TCP/IP. Many networks have reached capacity for their existing IBM 3745s. Instead of investing more money in older technology, organizations are migrating to multifunction channel solutions. For IP traffic using CLAW Packing or IP over MPC+ (which are recommended,) the throughput can be as high as 12 MB for bulk data transfer. For SNA using CSNA with bulk data transfer it can be upward of 8 MB. These speeds are for 1500-byte maximum transmission units (MTUs). In a 4096-byte MTU environment, the numbers can be higher.

In summary, the Cisco channel-attached routers offer many benefits such as speed, connectivity, and flexibility to an IBM enterprise network. By minimizing the number of FEPs required in a network, the channel-attached router solution offers a means to reduce network costs while improving performance.

Using an OSA-Express and Catalyst Switch Solution

The Cisco IOS Software provides the mainframe-required SNA and IP support described in the section Using Cisco Solutions as a FEP Alternative. These solutions can be placed at the core, aggregate, and branch level of the network. High-speed physical connectivity is the major differentiator when comparing the Gigabit Ethernet-attached solution to the ESCON-attached solution. The other differentiator is that the OSA-Express at its highest speeds supports only IP traffic.

IBM OSA

The IBM OSA is an integrated communications adapter for S/390, ESCON-based mainframes that provides direct attachment to Ethernet, Fast Ethernet, Token Ring, FDDI, and ATM networks. IBM previously released two versions of the OSA: OSA1, a two-card set that included two Intel 486 processors running channel offload software; and OSA2, a single-card replacement for the OSA1. The OSA2 eliminated the channel offload software required on the OSA1. The OSA-Express is the third iteration of OSA. With the availability of IBM's high-speed OSA-Express, customers must decide when to use the OSA-Express, and when to use the CIP or CPA. The following sections describe the OSA-Express, when to use it, when not to use it, and when the CIP or CPA is a better solution.

When to Use the OSA-Express

As IP-based applications on the mainframe require more and more bandwidth, the Gigabit Ethernet OSA-Express is becoming the method of choice for connecting the S/390 and z-Series mainframe to a TCP/IP network. The OSA-Express architecture removes the limitations of the channel protocols and places the S/390 on the same plane as the large UNIX servers. By using Queued Direct Input Output (QDIO) to the Self Timed Interface (STI) bus, the Gigabit Ethernet and Fast Ethernet cards have direct access to the 333-MBps (Generation 5 and 6 models) or 2-GBps (z/900 model) CPU buses. The OSA-Express is considerably faster than the current ESCON technology.

By rewriting the TCP/IP stack to use the DMA protocol against the OSA buffers, IBM has eliminated many of the buffer copies, which results in better throughput and reduced CPU resource consumption. IBM reduced the amount of configuration that is required for TCP/IP pass-through by loading the parameters from the TCP/IP profiles dataset, which eliminates the need to use the OS/2 or Windows-based OSA Support Facility (OSA/SF). The OSA-Express supports the Service Policy Server in S/390 and has four output queues, each of which is associated with a ToS. Prioritized by the Service Policy Server, application data is queued outbound. ToS bits are set and read by the Cisco network, providing end-to-end QoS. In general, use the OSA-Express for high-speed TCP/IP access connected to a Cisco network.

To benefit from the OSA-Express, a system must meet the following requirements:

- *IBM mainframe*—The mainframe must be an IBM mainframe. The OSA-Express is not supported on non-IBM hardware.
- *Generation 5 or later mainframe*—The OSA-Express is not supported on IBM processors prior to Generation 5. The Generation 5 mainframes became available in 1998.
- *OS/390 Version 2, Release 7 or later*—The operating system must be Version 2, Release 7 or later. QDIO, which enables the OSA-Express performance, was introduced in Version 2, Release 7.
- *IBM CS/390 TCP/IP stack*—The mainframe must be using the IBM TCP/IP stack.
- *TCP/IP access*—The QDIO and Direct Memory Access (DMA) improvements in the OSA-Express are for IP traffic only.

SNA traffic is transported over the high-speed OSA-Express as IP packets. The most common support is through SNASw or another EE implementation (APPN/HPR/IP). Layer 2 SNA traffic support is provided on the lower-speed OSA-Express cards, but they do not use QDIO and DMA and represent an expensive use of a valuable ESCON card cage slot. The OSA-Express is a high-speed channel into the mainframe. It is not a router or a switch and, therefore, must be connected to the network through a router or a switch. OSA-Express is a good choice for accessing the TN3270 server application on the mainframe, for high-speed FTP traffic, and for HPR/IP traffic.

When Not to Use the OSA-Express

Some customers cannot benefit from the OSA-Express because their environments do not meet the required criteria. Even when the customer does meet the basic hardware and software requirements, it may be more cost effective to use alternatives when SNA is transported natively. For example, in addition to supporting high-speed TCP/IP access through Gigabit Ethernet and Fast Ethernet interfaces, the OSA-Express supports lower-speed LAN interfaces, such as 4-Mbps Token Ring, 10-Mbps Ethernet, Fast Ethernet, and ATM. OSA-Express supports native SNA traffic in the same manner as the OSA2 card. When used in the OSA-Express, the lower-speed interface cards run at wire speed, which is not true of these same interface cards in the OSA2. Even at wire speed, OSA cards do not represent a good use of valuable mainframe real estate for SNA traffic.

The OSA cards are placed in a slot in the card cage on the mainframe. This card cage has a limited number of slots to connect a limited number of devices. A slot can be used for either ESCON or OSA connectivity. A single slot can support either four (pre-Generation 7) or 15 (Generation 7 and above) ESCON ports. A single slot can support four ESCON ports. The maximum theoretical throughput of these four ports is 68 MBps (4 x 17 MBps). If this same slot is used for a Gigabit Ethernet OSA-Express, the maximum theoretical throughput is close to 120 MB, which is a good tradeoff. If this same slot were used for a 10-Mbps Ethernet OSA card, the throughput would be less than 1.5 MBps, which is not a good tradeoff. The same reasoning applies to the Fast Ethernet versions of the OSA-Express. In these situations, a router should be used to aggregate WAN connections and to send aggregated data through a CIP or CPA. To enable SNA support, OSA/SF, an OS/2 and Microsoft Windows-based graphical configuration tool, must be used. (OSA/SF also has a 3270-style REXX interface.) The chief complaint of customers is that using the OSA/SF is cumbersome.

When to Use a CIP or CPA

Use a CIP or CPA in the following situations:

- *Non-IBM mainframes*—Use the CIP or CPA with any mainframe that supports the ESCON channel protocol, which is 100 percent of all IBM and IBM-compatible boxes. The OSA-Express is not an option for non-IBM mainframes.
- *Older mainframes prior to Generation 5*—Use the CIP or CPA on the approximately 60 percent of mainframes that do not support the OSA Express.
- *Older operating system releases prior to Version 2, Release 7*—Use the CIP or CPA with any currently supported operating system release.
- *Aggregation of TCP/IP and SNA traffic*—The CIP and CPA are an efficient use of the I/O card cage resources. Use the interface cards in the router to aggregate WAN traffic and to efficiently transport the combined traffic across the ESCON channel.
- *Offload processing*—Use the dedicated CPU and memory of the CIP or CPA 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 also be used to offset the inefficiencies associated with the mainframe TCP/IP stack in older releases (Version 2, Release 4 and earlier).

Why a Cisco Network Should Be Used with the OSA-Express

The OSA-Express provides high-speed TCP/IP access to the mainframe; however, it functions only as a network interface. The OSA-Express does not provide routing or switching, so it is dependent on the external network to perform routing and switching. This section describes the advantages of using a Cisco network in front of the OSA-Express. A Cisco network with a Catalyst 6500 Series Switch adds value to the OSA-Express in the following four major areas:

- Jumbo frame support
- End-to-end QoS
- Load balancing through Multinode Load Balancing (MNLB) features
- Redundancy

Jumbo Frame Support

A Jumbo Ethernet frame is 8992 bytes and is typically referred to as a 9-KB frame. A normal Ethernet frame is 1492 bytes and is typically referred to as 1500 bytes. Using the Jumbo Ethernet frame can result in significant increases in throughput (up to 75 percent) for bulk data transfers, such as file transfers or storage backups. To use Jumbo frames, all network devices between the servers must support the larger frame size. One benefit of using a Cisco switched network with the OSA-Express is that the Gigabit Ethernet interfaces on the Catalyst 6500 switches support Jumbo Ethernet frames. The Gigabit Ethernet interface card supported by the OSA-Express also supports Gigabit Ethernet interface cards on other IBM servers, such as the RS/6000, which also supports Jumbo Ethernet frames.

End-to-End QoS

The Service Policy System of the OS/390 operating system allows application traffic to be prioritized by application name, time of day, origin, destination address pairs, and so on. The Workload Manager uses this information to prioritize dispatching applications. The OSA-Express also uses this information to prioritize outbound traffic. The OSA-Express uses one inbound traffic queue and four outbound traffic queues. The outbound traffic queues are associated with the four TCP/IP ToS precedence settings. In conjunction with the Service Policy System, the S/390 can prioritize application performance, as well as outbound traffic.

Cisco and IBM have conducted interoperability tests. These tests prove that a Cisco network recognizes the prioritization of traffic from the OSA-Express and enforces this priority end-to-end through the network. The Cisco network uses features such as weighted fair queuing, custom queuing, priority queuing, and weighted random early detection to enforce the prioritization of the application traffic.

Note: In Cisco routers, weighted fair queuing is on by default on serial interfaces at speeds up to T1/E1 rates, which means that nothing extra is required to support the end-to-end prioritization of application traffic.

High Availability

One of the greatest differentiators for the IBM mainframe environment is its unrivalled availability. IBM mainframe environments are known for 99.99 percent availability, which equates to less than 15 minutes of downtime per year. If an application server has this type of uptime characteristic, it should be attached to a network with the same reliability. Testing conducted by IBM and Cisco has shown that one can design networks to eliminate any single point of failure in the network and the application server complex. During testing, which used both CIPs, CPAs, and OSA cards, the following scenarios were created and the application availability results noted:

- *CPU failure*—Application traffic was rerouted to another server in the complex
- *Logical partition (LPAR) failure*—Application traffic was rerouted to another server in the complex
- *Router failure*—Application traffic was rerouted through another router in the network
- *OSA failure*—Application traffic was rerouted through another OSA to the application VIPA address
- *CIP/CPA failure*—Application traffic was rerouted through another CIP/CPA to the application VIPA address

CONCLUSION

The IBM 3745/3746 FEP is the oldest, most mature IBM mainframe channel connectivity device. Groomed and enhanced over a period of several decades, the FEP has a great deal of functionality for SNA networks. Although there are some functions that neither a router attachment nor an OSA/E attachment can duplicate, the large majority of SNA and TCP/IP traffic can be very effectively and easily supported with them. Enterprises should evaluate their requirements carefully to see if they are even using any of the functions, and their goal should be to replace all local and remote FEPs. The rest of the SNA and TCP/IP traffic will be better served by gaining access to the mainframe via one of these alternate solutions.

APPENDIX A: FEP FUNCTIONS WITHOUT SIMILAR CISCO SOLUTION SUPPORT

A number of FEP functions cannot be replaced with a Cisco solution. These include the following:

- Multilink transmission group (MLTG)
- APPN (and LEN) Composite Network Node (CNN)
- XRF
- NRF
- X.25 SNA Interconnection (XI) and Network Supervisory Function (NSF)
- Emulation Program (EP)
- Dial Bisync 3280
- NTO
- NSI
- Message Entry and Routing with Interfaces to Various Applications (MERVA)
- Multisystem Networking Facility (MSNF)
- Teleprocessing Network Simulator (TPNS)

APPENDIX B: REFERENCES

For more information, please see the following resources:

- Cisco IBM Networking Protocols Website at http://www.cisco.com/en/US/tech/tk331/tk336/tsd_technology_support_protocol_home.html
- Cisco Business Ready Data Center Networking Solutions Website at <http://www.cisco.com/go/datacenter>
- Business Ready Branch Networking Solution at http://www.cisco.com/en/US/netsol/ns477/networking_solutions_packages_list.html
- Cisco/IBM Branch Transformation Website at http://www.cisco.com/en/US/partners/pr67/pr30/pr216/partners_strategic_solution_concept_home.html
- “Troubleshooting SNASw: A Guide to Problem Resolution” white paper at http://www.cisco.com/en/US/tech/tk331/tk336/technologies_white_paper09186a008011ff1b.shtml
- *SNA Switching Services Design and Implementation Guide* at http://www.cisco.com/application/pdf/en/us/guest/tech/tk897/c1501/ccmigration_09186a008007ca58.pdf
- “Token Ring-to-Ethernet Migration” business case at http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns74/ns149/net_business_benefit09186a00800c92b9.html
- “DLSw+ FID4 LLC2-to-SDLC Conversion for PU 4/5 Devices” tech note at http://www.cisco.com/en/US/tech/tk331/tk336/technologies_tech_note09186a0080093ec9.shtml
- TN3270 Server Protocol Website at http://www.cisco.com/en/US/tech/tk331/tk339/tk658/tsd_technology_support_sub-protocol_home.html
- RFC 2353 for a good general description of Enterprise Extender standards (see also IBM online documentation at www.ibm.com)
- Electronic Data Interchange-Internet Integration Group Applicability Statement 2 at <http://www3.ietf.org/proceedings/05nov/ediint.html>
- Cisco AVVID Website at <http://www.cisco.com/en/US/netsol/index.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.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 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 Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • 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 © 1992–2005 Cisco Systems, Inc. All rights reserved. Catalyst, Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0502R) DP/LW7991 05/05

