NCP and 3745/46 Today

IBM

Networking Controller Products, 3745 & 3746 News, Summer 2000

NTuneMON CD Inside

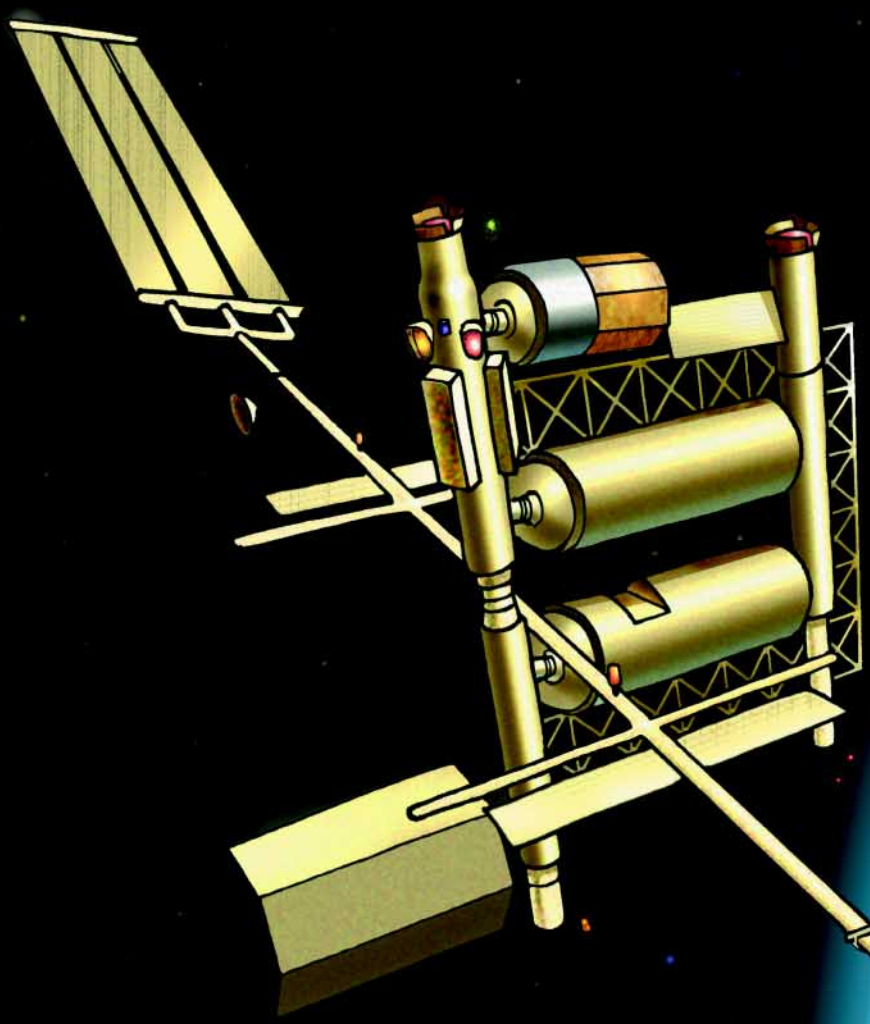

Evolving your SNA Investment to e-business

See page 12

IBM's Shark Enterprise Storage Server

See page 42

What's New with ACF/NCP and SSP? *See page 67*

Get More NTuneMON for Your Money *See page 69*

Contents

Continued...

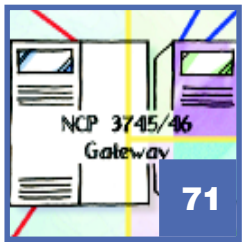

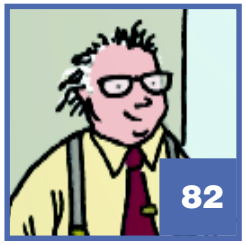

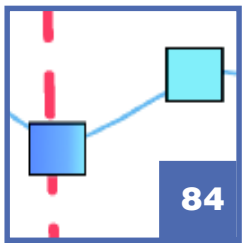

About This Issue

Once again, "Welcome" to *NCP and 3745/46 Today*. Well, "Y2K" is quickly becoming old news and most of our environments are back to an ever-quickening pace of installing and using new products, and providing and using new services. Change is what our business world is all about today, and so goes *NCP and 3745/46 Today*.

In this edition, we have taken a bolder turn away from traditional NCP and SNA networking topics with some interesting information about the "Shark" Enterprise Storage Server™—the IBM leadership statement in the disk-storage arena. Do you find yourself wondering why *NCP and 3745/46 Today* is publishing an article about disk storage? It's really not that strange when you consider who we have become and where storage and networking are going.

The largest majority of what used to be the Networking Hardware Division of IBM is now part of the Storage Subsystems Division (SSD) of IBM. NCP and 3745/46 products are part of the SSD. We've got a new name and some new associates, but we've got the same goal to provide you with quality products and services.

In addition, we have included a couple of strategy-statement articles about the task of migrating to IP networks. How do you identify when change is necessary? How do you identify what change is necessary? How can IBM help with the change? We've tried to address these questions, and more.

Please note that in the product section of this issue there are also some articles about new functions for NCP, SSP, NTuneMON™ and the 3746. I believe that you will find them interesting. Each new function is in response to a direct request from one of you, our customers. For those of you who are not already familiar with the many functions of NTuneMON, we have included an NTuneMON Demonstration CD in the back cover of this issue. This CD includes a brief overview of NTuneMON and a trial copy of the actual product for you to install and use in your OS/390® environment.

Thank you for being our customers and for the interest that you have shown in our products and our magazine. We are proud to tell you that our efforts to bring you an eye-opening, technically competent, magazine have already been recognized by the Society for Technical Communication (STC)—Carolina Chapter, which presented *NCP and 3745/46 Today* with a Merit Award at the 2000 Awards Banquet in March.

I am confident that after you have taken a look at this issue you will believe us when we say that we don't intend to rest on our laurels.

If we can answer any questions not covered here, please use the feedback form at the end of the magazine to contact us by mail or on the Web.

We hope you have a productive and successful year 2000.

John Woods

Manager of Networking
Controller Products

SNA and IP Integration: Is it Time to Move?

by Roger Messenger

In today's business environment, e-business and the Internet are known to the world's major corporations as a new and exciting way of conducting business. Intracompany and company-to-company transactions are being simplified, markets previously unreachable are being exploited, the cost of conducting business is being reduced, and investments are growing.

Traditionally, organizations carry out their day-to-day networking business using SNA applications, which are known for their robustness and suitability for managing mission-critical data. However, to exploit e-business and the Internet, the Internet Protocol (IP) is rapidly becoming the protocol of choice for the network backbone.

Cynics would say that SNA is dead. That is not the case—SNA still transports at least half of the world's business traffic. According to a survey conducted by CIMI Corporation[1], the amount of SNA-supported mission-critical application traffic is growing.

The survey further showed little to no movement of applications away from SNA. Many of the companies polled indicated that they would not begin to rewrite their SNA applications for at least 60 months. The author of the survey concluded that "while planners and managers know that, someday, their mission-critical applications will have to be changed, that day remains beyond their planning horizon."

Quality of Service and Security

The two major factors that make network managers and IT directors pause before migrating mission-critical applications from SNA to IP are quality of service (QoS) and security. It is QoS that gives SNA its unique ability to provide predictable response times to users throughout the network. QoS prevents interactive traffic from being overloaded by batch work. Data packets are not discarded, ensuring that the high availability and performance necessary to support your business needs are fully maintained.

The other factor—security of corporate data—is foremost in the mind of every IT executive. IP networks of today do not afford the protection that mission-critical applications need to maintain business integrity. Examples of the need for a high level of security are plentiful. Articles appear in the press on a regular basis reporting how corporate computer centers have been "hacked" into, compromising client files. Also noteworthy is the hackers' determination to disrupt network operations. Recent incidents, known as *Denial of Service* attacks, caused the networks of several large corporations to crash and be without service for several hours.

Convergence and Coexistence of SNA and IP

Corporations need to implement Web services and adopt an e-business infrastructure to compete with the small Web-based organizations that are being created almost on a daily basis to challenge established businesses for the e-business dollars. Internet applications demand the use of IP protocols. SNA and IP coexist to some extent in most major networks. However, for reasons of cost, efficiency, flexibility, manageability, and the pressure to compete on the Internet, network administrators recognize the need to begin converging their SNA and IP traffic onto a single infrastructure. The convergence of SNA and IP is further highlighted by the ever-expanding move to convert current business models to Web-based models.

These factors necessitate a network strategy that preserves the values of both SNA and IP, and that also contains costs and minimizes disruption to the evolving business model during the transition to a converged infrastructure.

At IBM, we understand the problems faced by many corporations that are implementing e-business and Web applications while, at the same time, depending upon SNA for mission-critical applications, and we have developed a strategy to help with this evolution. The foundation for this strategy was laid in 1999 when we enhanced all major components of the IBM 3746 Nways® Multiprotocol Controller Models 900 and 950 to meet the ever-growing need for improved performance and capacity.

[1] Nolle, Tom. "Will IP Replace SNA For Mission-Critical Applications?" *Business Communications Review*, October, 1999

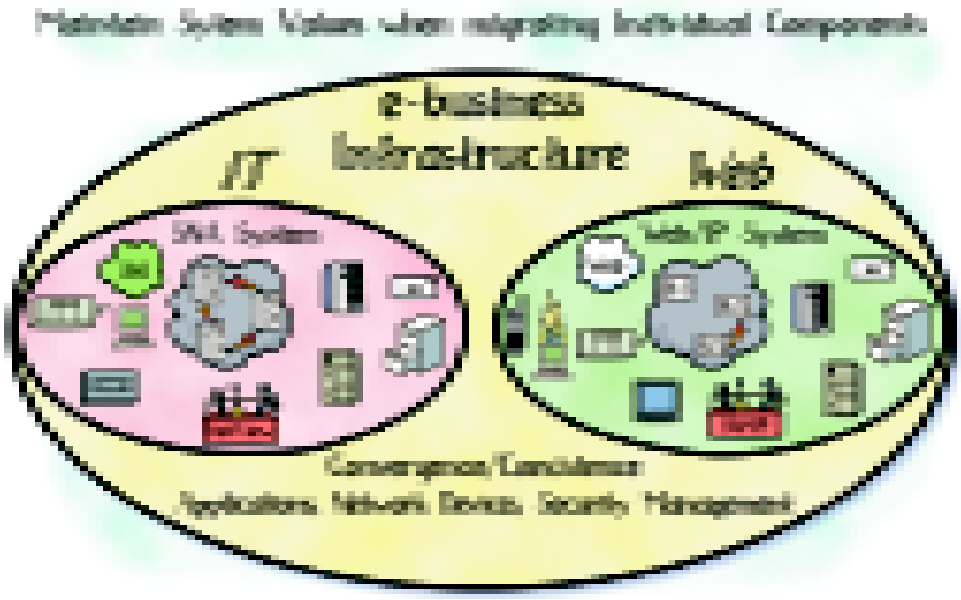

An article by John Foley (IBM 3746 Worldwide Brand Marketing Manager) detailing these enhancements appeared in the Fall 1999 edition of the *NCP and 3745/46 Today* magazine, and is recommended reading. To ensure that the needs of our SNA customers are fully met, we are continuing to announce additional enhancements to the 3746 family during 2000.

Technologies that Facilitate SNA-to-IP Convergence

The 3746 Nways Multiprotocol Controller is highly scalable and well positioned to accommodate future growth. In addition, this article describes several technologies that can be strategically deployed and that permit SNA data to be carried across an IP infrastructure, allowing you to capitalize on your 3745/46 investments.

- *TN3270E* is probably the most prevalent of these technologies. According to Tom Nolle (a nationally recognized consultant and Founder and CEO of CIMI Corporation), TN3270E has grown fivefold over the past 18 months. TN3270E is normally used in a subarea environment where the old "green" screens (or "dumb" terminals) in the remote environment have been replaced by programmable workstations such as PCs.

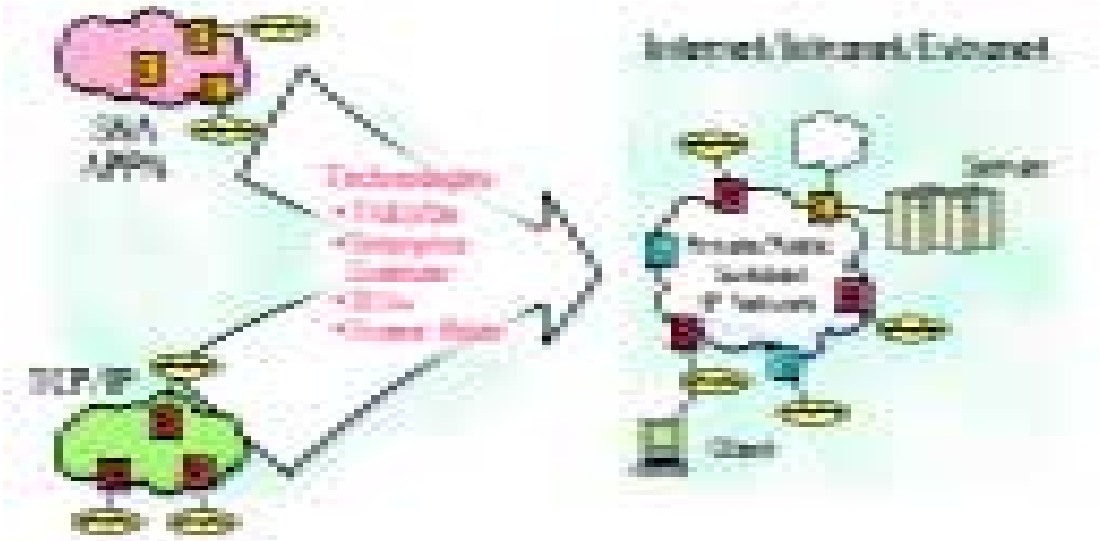

- *Enterprise Extender* is typically found in a S/390® Sysplex or Parallel Sysplex® environment where the business operation is dependent on high availability. To achieve minimum downtime, APPN® is implemented within the Parallel Sysplex system in the data center. To replicate high availability to the branch-office level, many organizations have extended APPN to the desktop.

  For those organizations deploying Parallel Sysplex and an IP backbone, Enterprise Extender is the ideal technology to support the transport of SNA (APPN/HPR) across the IP infrastructure, because the benefits of SNA, such as class of service (CoS) and non-disruptive route switching, are retained.

- *Data link switching (DLSw)* can carry both SNA subarea traffic and APPN. Note that DLSw does not support HPR and it does not provide CoS. DLSw is useful where it is not cost-effective to move end users from "dumb" terminals to programmable workstations. The downside of DLSw is that it creates single points of failure in the network and requires considerable processing power in the routers.

- *Frame relay* is a widely used transport protocol that enables IP, SNA , and other types of traffic to be carried on the same physical network. Typically, a frame-relay service from a public carrier is employed. Many users believe that frame relay provides better security than a private IP switched network. And, with the use of frame-relay boundary access node (BAN) for SNA applications, minimal setup is required and intermediate routers do not need large processors or massive amounts of memory to support the service.

OSA-Express

Supporting your move to e-business and Web-based applications is a main focus of the IBM S/390 Server Division. During the past three years, the cost of mainframe processing has fallen considerably, while performance has improved significantly. The release of OSA-Express in June 1999 made the S/390 server an excellent vehicle to manage e-business and Web applications. When a business adopts an e-business strategy, its network and server infrastructures need to support an exponential growth in data flow.

The complex graphical design of Web sites means that data transfer is reckoned in terms of gigabits rather than the kilobits and megabits of previous applications. For an organization to maintain a successful Web base, it must be able to carry data in and out of its server at megabyte speeds. Today's consumers are fickle—if a target Web site proves difficult to access, or response times are considered excessive, the end user will quickly switch to a better-performing site, never to return.

OSA-Express with its support for Gigabit Ethernet, attached to a Layer 3 switch, will become the S/390 access of choice for many organizations as they build complex Web-based applications. To support the high data-throughput requirements of OSA-Express, the S/390 Server Division has made some important architectural changes to the G5 and G6 range of S/390 servers. OSA-Express has a direct attachment to the Self-Timed Interconnect (STI) bus, which runs at 333 MBps. A technique called *Queued Direct Input/Output (QDIO)* enhances the data path between OSA-Express and the S/390 memory. It minimizes I/O interrupts, thereby reducing the load on the processor. Furthermore, it reduces TCP/IP path lengths by 25%.

TCP/IP data carried through OSA-Express has been measured in and out of the S/390 server in full-duplex mode at 50 MBps. This is a 5:1 performance improvement when compared with an ESCON® channel running at 10 MBps. Cost savings over channel-attached routers are also significant.

Current releases of the OS/390® software support technologies that allow SNA application data to be carried over an IP infrastructure. For example, for those organizations migrating from a subarea environment, the S/390 can be used as a TN3270E server. It supports over 64 000 sessions, thereby eliminating the need for costly external TN3270E server devices.

In the Parallel Sysplex with APPN/HPR, where high availability is paramount, Enterprise Extender can be terminated in the S/390 server. Doing so allows APPN/HPR to be carried directly to the application, which retains the advantages of SNA transport, but over IP.

For those data centers that are not ready to manage gigabit speeds, OSA-Express supports auto-sensing Fast Ethernet and ATM OC-3 interfaces.

OS/390 software now provides Policy Management, which is important to enterprises that are focusing on e-business and Web-based applications. Policy Management allows Type of Service (ToS) settings to be used, which enforce an IP QoS and bring more predictable response times to the end user.

Many data centers are not at the right level of hardware or software to fully support OSA-Express or the enhanced TCP/IP stack. A minimum requirement for OSA-Express is S/390 G5 and later, coupled with OS/390 Version 2 Release 7 for Gigabit Ethernet support. Because of hardware and software prerequisites, many locations might not reach the appropriate levels for OSA-Express implementation within the next 18 months. In some environments, the process might not start until much later, so it is assumed that channel connectivity will remain in data centers for several years. For this reason, it may be necessary to leverage installed 37xx equipment with hardware and software upgrades until the next step to OSA technology can be implemented.

For a more detailed discussion of OSA-Express and its capabilities, see the article "It's Time to Switch—OSA-Express Drives the TCP/IP Data Center and the Network Together" by Jim Goethals on page 10 of this issue of *NCP and 3745/46 Today*.

The figure shows that OSA-Express will be the technology of choice for accessing TCP/IP in the S/390 in the future. Recognizing the importance of OSA-Express in the longer term helps us to understand the evolution options for 3745 and 3746. Given that the preferred option for users with the correct versions of S/390 hardware and OS/390 software will be to migrate to OSA-Express in the medium-to-long term, the short-to-medium term solution now becomes clearer. The 3745 and 3746 platforms are ideally placed to provide solutions for all the SNA requirements and early TCP/IP environments in the medium term (2 to 4 years).

There might be a temptation to move to channel-attached routers. But you should stop and consider whether this is strategically wise as you look at your long-term plans for S/390 access. Migrating to channel-attached router technology in the near term could be an expensive and unnecessary step. Routing technology has a limited life span, whereas switching technologies provide greater speeds and efficiencies. S/390 connectivity in the future will be achieved through switching into OSA-Express. Consequently, the network gateway will be a Layer 3 LAN switch, not a router. This means that migration to a channel-attached router solution at this stage could involve two forklift upgrades, and is unlikely to yield a good return on the investment.

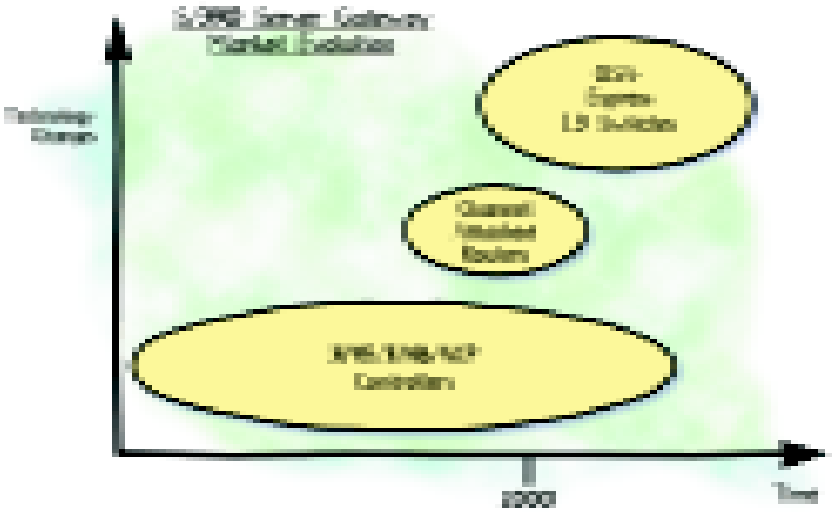

Roger Messenger (roger_messenger@uk.ibm.com) is a member of the World Wide Brand Marketing team and is based in the U.K.

It's Time to Switch—OSA-Express Drives the TCP/IP Data Center and the Network Together

by Jim Goethals

OSA-Express—an IBM-exclusive innovation on the S/390 Parallel Enterprise Server™ G5 and G6—makes the direct connection to your IP network switch, a data-center design imperative. This third-generation Open Systems Adapter (OSA) is the most cost-effective, highest-performing, and best price-performance S/390® network connection to date. OSA-Express connects the power and performance of the S/390 to your ever-growing e-business transaction workload. OSA-Express is a key element in optimizing S/390 TCP/IP access, providing the basis for Web serving as well as SNA-to-IP investment evolution.

A New Adapter Technology—OSA-Express

OSA-Express offers full-duplex, direct attachment to the S/390 333-MBps Self-Timed Interconnect (STI) bus. This is a significantly faster attachment than ESCON® networking devices or OSA-2 connected to the 17-MBps Channel Request Handler bus.

OSA-Express consists of three different hardware features. All of them have a single port, which can be shared among logical partitions (LPs) by using the Enterprise Multiple Image Facility. A maximum of 12 OSA-Express adapters can be connected to a S/390 Parallel Enterprise Server.

The OSA-Express features are:

1. The *Gigabit Ethernet (GbE) feature*, which operates at 1 Gbps and includes support for up to 9000-byte jumbo frames. This support is especially helpful for file transfers and large data transfers such as backup recovery with Tivoli® Storage Manager. The OSA-Express GbE supports TCP/IP protocols only. SNA applications can communicate through the OSA-Express GbE adapter by using SNA-to-IP integration technologies such as TN3270E servers, WebSphere™ Enterprise Edition, and the Enterprise Extender function on the S/390. Enterprise Extender provides APPN®/HPR over the TCP connection. This adapter has been available since June 25, 1999.

2. The *Fast Ethernet (FENET) feature*, which can auto-sense both 10-Mbps and 100-Mbps speeds, while providing protocol support equivalent to the prior OSA-2, but now at wire speeds. This includes TCP/IP, SNA subarea, APPN/ISR and APPN/HPR support. This adapter has been available since January 31, 2000.

3. The *ATM feature*, which operates at 155 Mbps while providing protocol support equivalent to the prior OSA-2, but now at wire speed. This includes ATM LAN emulation, and ATM native support for both TCP/IP (Classical IP or RFC 1577) and APPN/HPR. LAN emulation supports both Ethernet and token ring. This adapter has been available since January 31, 2000.

A bonus feature for users of the OSA-Express Ethernet environment is called *LP to LP Communications*. This function provides 333-MBps data transfer between logical partitions connected to (sharing) the same OSA-Express adapter. The data is sent directly from one logical partition to another without traversing the LAN. Supported adapters include Gigabit and Fast Ethernet, and Ethernet LANE on the 155 ATM adapter. This feature gives an added boost for those applications that require communication across logical partitions.

A Significant Enhancement—QDIO

A companion enhancement specific to OSA-Express is *Queued Direct Input/Output (QDIO)* architecture. QDIO, which supports IP packets in the Ethernet environment, is a significant enhancement to the I/O subsystem of OS/390® Version 2 Release 7. QDIO transfers data from the adapter's memory *directly into S/390 memory*, bypassing three layers of processing inherent in the prior Channel Request Handler technology used by ESCON devices and OSA-2.

This means that a QDIO-defined OSA-Express provides continuous data exchanges using memory queuing with minimal I/O interrupts and continuously active communications. The IP Assist function reduces S/390 cycles required for compute-intensive networking functions such as Address Resolution Protocol (ARP), MAC address handling, and more. The new S/390 Signal Adapter instruction notifies OSA-Express when outbound data is ready for transmission, a significant savings versus prior methods. And by using direct memory access (DMA), QDIO employs an advanced Programmed Controlled Interrupt Algorithm, which allows OSA-Express to minimize S/390 interrupts, significantly reducing CPU utilization.

Configuration and administration of the adapter is significantly simplified with the dynamic OSA Address Table. The definition of IP addresses is done in one place, CS OS/390. The information is passed to the adapter at IP device start. The OSA Address Table is built and updated dynamically, simplifing installation and configuration setup, and eliminating data-entry errors and incompatible tables. Optimized for the TCP/IP-based world of e-business, QDIO reduces TCP/IP path length by 25%.

QDIO is currently available on the OSA-Express Gigabit Ethernet adapter. Support for QDIO on the Fast Ethernet and the Ethernet LAN emulation of 155 ATM adapters has been announced for planned delivery 3Q2000. This means that the Fast Ethernet adapter and the 155 ATM adapter used for LANE will be capable of operating in either of two, mutually exclusive operational modes: QDIO for TCP/IP in Ethernet environments only; and non-QDIO for support equivalent to OSA-2 including TCP/IP (using LAN Channel Station Protocol), and SNA, APPN ISR and HPR support (using Link Services Architecture Protocol).

This Combination Demands Attention

The price-performance advantages of OSA-Express are compelling. As illustrated in the figure, just compare what it requires to bring 50 MBps of bandwidth into the S/390 for the backup of a file:

- One OSA-Express Gigabit Ethernet adapter with a single attachment to a Gigabit Ethernet switch requires one IP address and one channel path ID defined.
- An ESCON channel-attached router with a minimum of five 10-MBps ESCON channel adapters and connections, a powerful router platform, a Gigabit Ethernet adapter to attach to the switch, five IP addresses and five channel path IDs to provide equivalent capacity.

The OSA-Express Gigabit Ethernet adapter provides *five times* the performance and *thirty-two times* better price-performance than an ESCON channel attached router for IP Gigabit Ethernet attachment!

For TCP/IP access to S/390 . . . it's time to switch to OSA-Express!

For more information about OSA-Express, go to: **ibm.com**/s390/networking/

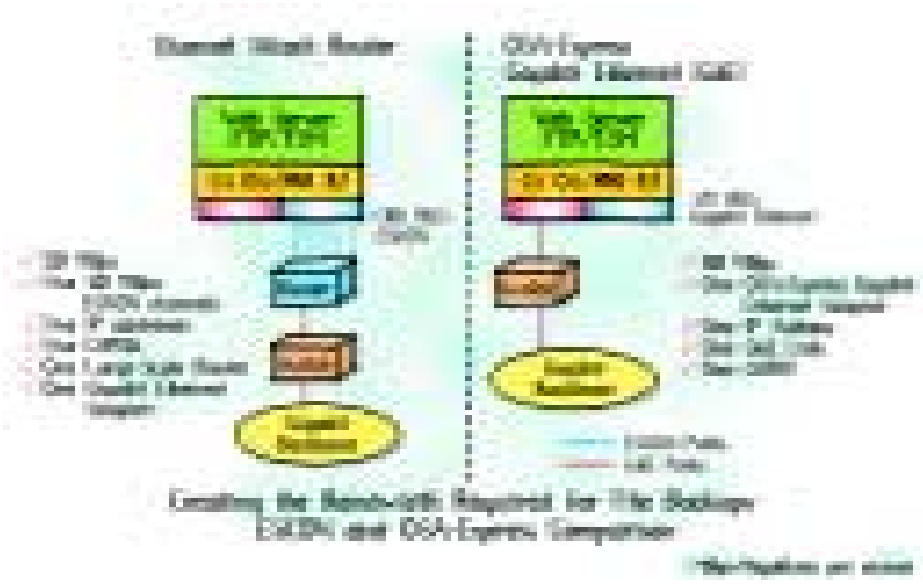

Jim Goethals (jimgo@us.ibm.com) has been associated with networking products 23 of his 31 years at IBM. He was most recently involved in product marketing of the IBM routing family including the 2210, 2216, the 3746 Multiaccess Enclosure, and the Network Utility. He is now the S/390 Networking Offering Manager. Jim specializes in large system network requirements and is a regular contributor in Enterprise customer briefing sessions conducted in the Raleigh International Networking Center.

Evolving your SNA Investment to e-business

by Richard Tobacco

"How can I get to the new e-business paradigm quickly without destroying all of my past investments?" CIOs everywhere have been asking.

"I have this extremely large investment in SNA applications, networks, and clients. The Internet is based on IP. How can I migrate these investments into an IT environment that fully leverages the burgeoning e-business marketplace?"

"Is the migration to e-business really such a big deal?"

To answer these questions, let's look at a few facts. First of all, most of today's business data continues to reside on a S/390® or equivalent server. In fact, in many businesses, almost every online transaction runs through a S/390 server or equivalent at some point. This data is hosted by mainframes running applications that were written to share data over an SNA network.

Secondly, it is prohibitively expensive either to migrate the data from these SNA-based applications to new IP-based applications or to modify the SNA-based applications so that they also support IP natively. Either option also tends to be fraught with other issues like availability of programmer skills and the migration complexities that surround mission-critical data.

Lastly, the growth in IP-based networks is increasing rapidly while the use of SNA networks and resources is decreasing at a significantly slower rate. This indicates that the aggressive leveraging of IP to move to new e-business opportunities has not produced a corresponding aggressive approach to moving away from SNA-based solutions.

Stages of e-business-Driven SNA Migration

Companies that have significant SNA investments are at various stages of leveraging the e-business revolution. Often, large corporations have groups with different business pressures in different stages of migration. We will simplify and refer to these business groups as companies. Companies generally fall into three stages of migration to e-business: *traditional*, *transition*, and *growth*.

- Traditional companies continue to maintain an exclusive SNA environment. They are driven by valid, absolute business demands like highly predictable performance, continuous availability for their applications and a very high level of scalability. These companies are simply not motivated to change this portion of their IT infrastructure right now.

- Transition companies are motivated by the need to reduce their infrastructure costs. However, they do not yet have a prevailing need to replace their SNA-based clients. These devices could be banking terminals requiring a particular SNA communication protocol (for example, SDLC) or they could be actual 3270-like devices, coax-attached to a 3274-like terminal controller. These customers are in the process of removing their SNA intranets while still maintaining a SNA presence on their back-end servers at the data center and on their clients in the branch office.

- Growth companies, while still having applications that leverage back-end SNA data, require an end-to-end IP presence in their intranets. They are motivated both by the need to pursue new e-business-related opportunities as well as the need to continue to reduce IT costs. These companies are making everything *look* like the Internet by using IP clients (for example, Web browsers and TN3270), providing clients access to their S/390-hosted data, and building *new* IP-based applications that work with their *existing* SNA-based applications. These companies are inventing new and creative ways for their employees, customers and partners to access their wealth of back-end data through standardized IP-based clients.

In 1998, companies were equally distributed across these three stages. However, industry projections reveal that about 75% of companies will be in the growth stage by 2002. In fact, there will soon be more IP-based business clients, networks and applications than their SNA equivalents. Because companies are starting from very different histories and investment strategies, there are valid business reasons for each IT shop being in a different migration stage. As a result, *IBM's commitment and capability in this revolution is to facilitate each company's e-business goals regardless of its SNA migration stage*.

IBM's product and service offerings today do just that. While we will address offerings appropriate for each migration stage, let us first remind ourselves that the e-business model assumes there is universal access to back-end applications via IP networks.

Enterprise Extender Advantages

- *Provides SNA data-path information to the network routers, virtually eliminating router overhead.*
- *Supports an IP connection into S/390® for the encapsulated SNA traffic—provides Gigabit Ethernet, Fast Ethernet and ATM connectivity.*
- *Eliminates the termination and acknowledgment workload, enabling routers to handle a much larger number of users.*
- *Provides failure-resistant access: supports the IP-reroute capability, maintains connectivity, and switches to an alternate path without session disruption.*
- *Provides port-level traffic priority to the encapsulated session data.*

Supporting Companies in the Traditional Stage

We will continue to support our existing subarea and APPN®/HPR offerings for the traditional company. Enhancements to APPN/HPR will continue, especially on the S/390 server where companies' expectations surrounding availability, redundancy, and scalability continue to grow. In comparison, traditional subarea support requires little enhancement.

In addition to our well-known 3746 Nways® Multiprotocol Controller models, we have full SNA APPN/HPR support in our software solutions. IBM Communications Server for OS/390®, OS/400®, AIX®, Microsoft® Windows NT® and OS/2® provides support for those customers maintaining their SNA networks.

Supporting Companies in the Transition Stage

Companies in transition are converting their SNA backbone to IP while maintaining their SNA-based servers and clients. These companies need to encapsulate their SNA-based traffic so that it can traverse the IP backbone. IBM created two open encapsulation technologies for doing this: data link switching (DLSw) and the newer, easier-to-deploy and more powerful Enterprise Extender (EE). While the EE technology provides multiple advantages for a company, investment history and technology skills will again be key decision factors.

DLSw is the encapsulation of SNA protocols within a TCP/IP "package." The package is a usable TCP/IP address that has been discovered by DLSw-capable routers. At both the sending and receiving location, the DLSw-capable routers terminate the SNA connection and locally acknowledge and control transmissions. This can be a significant burden even to powerful routers, and DLSw solutions often require multiple DLSw-capable devices. Because DLSw was introduced as an open standard, many businesses rely on this encapsulation scheme.

Evolution of SNA Transport

Original versions of SNA, designed to support networks with unreliable links, were very tightly coupled with data integrity that was checked at each path step. Improved link quality enabled SNA to be restructured with data verification only at the endpoints. This advanced SNA with High-Performance Routing (HPR) separated transport functions from data-integrity functions.

The similarity between HPR—an improved Advanced Peer-to-Peer Networking® (APPN®) routing protocol—and TCP/IP might be apparent:

- *End stations ensure data integrity and intermediate devices forward traffic.*
- *The intermediate node forwards packets, has no session awareness, and relies on the endpoints for error recovery.*
- *The endpoint nodes provide end-to-end error recovery, non-disruptive rerouting and selective retransmission of lost packets.*

A difference between SNA and TCP/IP transport is that HPR sessions use the same path as long as it is available, whereas TCP/IP sessions might change paths. Because the same path is guaranteed, the intermediate SNA routing devices rely on the endpoints for a routing-path update. This eliminates the need for the routing devices to maintain the session route information.

Host On-Demand

A Java™ application provides a TN3270E client that can be downloaded from a server. This server can be different from the application host. Host On-Demand is a 2-tier solution that removes the server from the data path. Application distribution and update is automatic, significantly reducing administration.

EE is an extension to High-Performance Routing (HPR), a lower-overhead, less-complex SNA, that provides efficient encapsulation of SNA application traffic within UDP/IP frames by EE-capable devices at the edges of an IP network. The SNA traffic is routed without hardware or software changes to the IP backbone. To the end-user, the session is "normal SNA" with predictable performance and high availability.

We provide EE support in our routers and Communications Server software for OS/390, OS/400, AIX, Windows NT and OS/2. All provide EE support for those companies migrating their SNA transport networks to IP. Like DLSw, EE is an open technology, and recently Cisco Systems, Inc. announced their routers would also provide the EE function under the name *SNA Switching Services (SNASw)*.

Supporting Companies in the Growth Stage

The basic strategy of companies in the e-business growth stage is to translate data to and from their SNA-based applications at the data center before it hits the WAN. There are essentially two ways for IP-clients to access SNA data today: Telnet clients and Web browsers. IBM provides these mechanisms for accomplishing this server-side translation.

Businesses can choose to convert all their end-user "terminals" to be TCP/IP-based like IBM Host On-Demand (HOD) or Personal Communications clients. Both of these TN3270 clients are available on multiple platforms—HOD on OS/400, AIX, and Windows platforms, PCOMM on Windows and OS/2. These clients are able to access Web-based applications, applications written for TCP/IP, and existing SNA applications with terminal emulation. To support these clients, IBM provides a robust set of Telnet servers across OS/390, OS/400, Windows NT, and AIX systems. IBM and Cisco routers also provide this TN3270 server capability. All these server applications translate SNA to IP and present the data to IP-based clients. Today, more businesses use TN3270 than Web-browsers as their method of accessing those SNA applications.

Host Publisher

- *Translation of back-end data into an HTML data stream provides universal Web-browser access. Host Publisher eliminates the dependency of clients on the correct software level for back-end data access.*

The expected predominant access technique in the near future involves Web browser clients, and IBM provides server-side Web-to-host integration solutions. This translation is provided across all of the above-mentioned environments. In addition, with its Host Publisher offering, IBM provides a very effective tool set for the growth customer to quickly create new IP-based applications. Leveraging back-end host data without having to change those back-end applications facilitates their race to e-business!

IBM's Host Publisher products are available on all of the platforms mentioned above—OS/390, OS/400, Windows NT, and AIX.

You Must Never Forget Management!

A significant part of a company's investment in SNA is its SNA-based management scripts and automation. It is critical to consider management systems as part of a company's overall migration plan. While some of these management processes are company-specific in that companies have written their own tailored automation scripts, the good news is that Tivoli® products manage pure SNA, pure IP, and SNA-over-IP scenarios.

What Migration Stage Are You in?

What stage of e-business-driven SNA migrations are you in? Do you know how to make your way to the e-business model? As your company plans its path to e-business success, remember that IBM provides you with many options today! You can maintain SNA in the data center for S/390 Parallel Sysplex® availability and non-disruptive growth. You can separately transition your SNA backbone to IP without having to displace your SNA clients. And you can leverage your IP clients, enabling a new class of e-business client with browser access to SNA applications.

Richard Tobacco (rjt@us.ibm.com) has been associated with sales support of networking products since the mid-80s. He is currently Brand Manager for Communications Server for OS/390, a product that provides TCP/IP and SNA connectivity.

3745/3746 Server Access Scenarios with Technology Overviews

by John R. Foley, Jr.

This article[1] identifies the placement of 3745/3746 controllers in classic SNA and APPN® networks and identifies the future role of the 3745/3746 controllers as they accommodate smooth transitions to emerging technologies such as high-speed Layer 3 switching. In addition, we review several key technology options that help you make informed decisions when merging classic SNA solutions with emerging e-business applications over internets.

3745 and 3746 (374x) Family Role in e-business Infrastructures

The 374x family is an essential component of the road map for the transition of the IT infrastructure toward full e-business enablement. IT infrastructures must incorporate both networks based on Web applications and networks running legacy applications. The role of the 374x family in the integrated IT infrastructure is shifting from that of network owner to that of providing access to mission-critical S/390® applications, as well as network support and connection to legacy devices whose continuing value does not justify replacement.

NCP provides unsurpassed value in enabling access to legacy devices and in complementing the VTAM® software on the S/390 server. NCP is an essential partner with VTAM in subarea environments. The NCP routes subarea traffic to the correct VTAM application server, eliminating the costs and delays that would be incurred if VTAM performed the routing. The NCP offloads SNA boundary function (session management and control) for the thousands of attaching devices from VTAM, thus reducing cost. Subarea environments without NCP may realize a notable increase in VTAM resources for the SNA boundary function and VTAM-to-VTAM routing. NCP and VTAM, as complementary partners, have together forged the IT infrastructure on which most of the world's businesses have been based.

The 374x family has extended its partnership with VTAM from the subarea environments to the advanced technologies of APPN and HPR, which enhance the S/390 Parallel Sysplex® configuration. In addition, the 374x family provides a second critical role in connecting the S/390 Parallel Sysplex with legacy and advanced SNA networks and devices. The 3746-900 with NNP and the NCP operate together to provide both subarea and APPN/HPR functions on the same platform. The third role of the 374x family is to enable the composite SNA and IP IT infrastructure for e-business by enabling S/390 access from both IP and SNA networks.

Consequently, the Parallel Sysplex®, when supporting IP traffic today, provides many of the innovations previously relevant only to the SNA environment. The 3746-9x0 with NNP has integrated IP routing and connectivity to provide access to S/390 TCP applications from IP networks. For those environments where the 3745 and 3746 carry both SNA and emerging protocols, there are various options. These are discussed in detail in this article.

[1]The information in this article has been excerpted from the *"S/390 Server Access Solutions—3745/3746 Evolution Strategy"* white paper authored by John Foley, Rachel Pickering, and Pratt Parish. This white paper in its entirety can be found at **ibm.com**/networking/nethard.html.

During 1999, major enhancements were announced on the 3746, providing increased performance and connectivity and enhanced management tools. Additional enhancements planned for the second half of 2000 include increasing the number of:

- LU-LU data sessions established by the 3746 Network Node
- SSCP-LU control sessions established by the 3746 Network Node
- PUs per Token-Ring Processor Type 3
- IP routes
- OSPF routes

In addition, the ESCON® Type 3 Processor is scheduled to be enhanced to provide additional IP throughput, and new RUNCMD commands will be provided to manage the 3746 from the Tivoli® NetView® platform. Concurrent access by multiple operators to the 3745/3746 through Java™ software consoles is also planned in 2000, as is providing for the automatic discovery of low entry networking (LEN) nodes. The latter will eliminate the need to predefine these network resources. Finally, IBM will announce a Service Update Package (SUP) for NCP Version 7 Release 8 that supports the 3746 enhancements along with a new release of NTuneMON™, which provides new tuning and display capabilities for the controllers' resources.

The main reasons why customers continue to depend on the solutions provided by the SNA protocols and networks are:

- Numerous mission-critical applications are VTAM/SNA-based
- SNA delivers high availability and is eminently scalable
- SNA has first-class management facilities, something that is essential to mission-critical business applications

For these reasons, the 3745 and 3746 family of products continues to be the platform of choice for access to S/390 servers and their mission-critical applications. The 37xx family provides a dependable infrastructure to support mission-critical IT applications with a full range of protocols (SNA subarea, APPN/HPR and TCP/IP), scalable connectivity, and the unique capabilities provided by the NCP and associated NCP Extended Options (NEO) licensed programs (X.25 NPSI, EP, NTO, NRF, NSI, XI, ALCI).

Technology Overview

As the world embraces the Internet and accepts e-business as the way to perform business, it's imperative to recognize that the majority of mission-critical applications residing on mainframes are still running SNA protocols. The major challenge for the network manager is to provide a means of integrating current SNA mission-critical traffic over backbones that are becoming predominately IP based. The real task will be to integrate these two distinctly unique protocols without causing disruption to network operations, compromising performance and most importantly, without introducing breaches in security.

Before exploring the methods in which you can capitalize on your current 37xx investment and use this platform to offer the smoothest transition to e-business, it's important to understand the technologies that are available today, which will help you transcend your network to a lean, mean SNA-IP solutions machine.

The following section discusses methods of transporting SNA traffic over an IP backbone and associated OS/390® operating system-based features that facilitate access to host applications.

Data Link Switching

Data Link Switching (DLSw) is one of the most common methods used today to transport SNA traffic over an IP-backbone network. Although initially developed by IBM, it is an open standard used by many manufacturers, and is based on a number of IETF standards, including RFC1434, RFC1795 and most recently, RFC2166. It uses a TCP connection to carry the SNA over the IP backbone; TCP provides the reliability and guaranteed delivery that is required for SNA applications.

DLSw defines and locates the partner routers, and sets up the TCP connection. The SNA layers at the session endpoints are unaware of the IP backbone and DLSw. DLSw routers are attached to the SNA network using LAN connections at each end. SNA traffic depends on link-level acknowledgments, and DLSw is responsible for providing these acknowledgments (known as *local LLC termination*). This avoids unnecessary LLC (LAN) keep-alive messages from flowing over the WAN.

The TCP connection that carries the SNA is set up between two routers, at the ingress and egress of the IP backbone. Although TCP carries the SNA reliably over the IP backbone, the two endpoint routers are single points of failure, and so DLSw cannot offer a completely reliable end-to-end service.

DLSw carries two types of SNA traffic: subarea SNA and APPN SNA. It cannot carry SNA HPR traffic. The SNA class of service (CoS) priority cannot be carried across the IP backbone, although techniques such as bandwidth reservation can be used on the first hop into the IP network.

DNS with Workload Manager Support

The Domain Name Server/Workload Manager (DNS/WLM) in Communications Server for OS/390 V2R5 allows for enhanced Telnet and FTP workload distribution in a Parallel Sysplex environment by factoring in the current load of each server. The new DNS takes input from the OS/390 WLM component, allowing for timely decisions to be taken about which IP address to use for a connection request. However, two major limitations exist in the DNS/WLM approach:

- The dependency in the DNS lookup at the S/390 each time a new connection is made
- The dependency on the clients not caching the IP address

Because many clients do cache IP addresses, the effectiveness of DNS as a basis for real-time load balancing might be low. As a result of this limitation, and the need for a more effective technology, IBM created Network Dispatcher.

Sysplex Distributor

The Sysplex Distributor is a new function in OS/390 Version 2 Release 10. It offers a number of benefits that address some of the configuration restrictions existing today with TCP/IP workload balancing and high-availability solutions. The Sysplex Distributor:

- Provides workload balancing without the use of the DNS/WLM or Network Dispatcher, supporting the customer's investment in Parallel Sysplex.
- Enhances VIPA take-over support, providing non-disruptive give-back of VIPA to the primary owner. Today, give-back to a VIPA primary owner is disruptive to existing TCP connections.
- Promotes use of OSA-Express, removing the Network Dispatcher (and Cisco Network Director) configuration restrictions of single-hop route.
- Promotes the S/390 server as Web server, removing the configuration restrictions of Network Dispatcher (and Cisco Network Director) and dependency on specific WAN hardware for workload balancing.

The S/390 Open Systems Adapter for Direct LAN Connection

IBM introduced the S/390 Open Systems Adapter (OSA) in 1995 as the strategic S/390 connectivity interface for networking. The OSA, OSA-2, and OSA-Express provide open, industry-standard LAN interfaces integrated in the S/390 servers to reduce S/390 total costs of computing and to simplify access to enterprise networks. OSA provides many benefits: it supports server-to-server communications, expands connectivity alternatives, and extends the S/390 strengths of security, backup/recovery, availability and scalability to the edge of the network.

Open System Adapter is unique among the S/390 access alternatives discussed here in that it is integrated within the S/390 server platform. The S/390 server places function among the various S/390 components. For example, the Enterprise Extender is supported on S/390 servers and is placed in Communications Server for OS/390 software with the protocol stack components, but it is not placed on the OSA itself. The S/390 server has rich networking functionality with placement of function optimized for the enterprise server environment. Therefore, some of the functional comparisons include features provided by other components of the S/390 server solution.

The OSA-Express configured with fast LAN technology, such as Gigabit Ethernet, delivers extraordinary growth in S/390 access capacity. For example, a single ESCON channel can achieve somewhat above 10 MBps throughput. OSA-Express configured with Gigabit Ethernet has a goal to exceed 100 MBps, which is a tenfold increase in access-link capacity. Environments that are OSA-Express-capable will reap the benefits of OSA-Express for IP.

Both OSA-2 and OSA-Express provide transport of TCP/IP, SNA (subarea, APPN and HPR) and IPX traffic to S/390 servers. However, the OSA-Express Gigabit Ethernet feature supports only TCP/IP.

OSA-2 and OSA-Express can be used in conjunction with the 374x, replacing the ESCON connection with an upstream token-ring connection. However, the link connectivity between the 374x and S/390 is not the primary consideration in many configurations. It is rather that of maintaining the cooperative partnerships between VTAM and NCP in subarea environments, and between VTAM (DLUS) and a DLUR such as the 3746 NNP. OSA alone does not include NCP or DLUR functionality. For a more detailed discussion of the OSA-Express, its prerequisites, and the functionality it brings, see the article by Jim Goethals on page10 of this edition of *NCP and 3745/476 Today*.

APPN and HPR

Advanced Peer-to-Peer Networking® (APPN) and High-Performance Routing (HPR) are open, industry-standard architectures published by the AIW[2] and implemented by IBM and other vendors. APPN and HPR are important in developing the IT infrastructure for e-business and are therefore significant functions of the 374x. Most mission-critical business information comes from an SNA heritage and even today, a large percentage of such traffic is based on SNA applications. The S/390 Parallel Sysplex has incorporated dynamics, scalability and workload-balancing functions using APPN and HPR. To efficiently access SNA data in this environment, APPN/HPR should be used inside the Parallel Sysplex and at least on the network gateway. Therefore, there is a need to establish an APPN/HPR strategy for the S/390 servers in the Parallel Sysplex configuration.

An important function of the 374x family continues to be that of enabling efficient, robust network access to S/390 applications. When S/390 servers are represented as APPN network nodes or APPN end nodes, the 374x family provides several gateway migration solutions to interconnect the SNA subarea network with the S/390 servers. The 374x network gateway allows the connection of an SNA subarea network to S/390 servers that can be running APPN/HPR, without the requirement for any network changes.

S/390 APPN Migration Options

There are a number of options when considering APPN migration in the data center, and many books have been written on this subject. Some of the relevant migration alternatives that are referred to in this article are:

- Migration data host (MDH): VTAM has both subarea and APPN identities and converts subarea traffic to APPN/HPR within the data center.
- Composite network node (CNN): NCP and VTAM (in partnership) give the appearance of a single APPN network.
- Dependent LU requester/server (DLUR/S): VTAM, in conjunction with an APPN network device, such as the 3746 Network Node Processor (NNP), perform a conversion so that subarea traffic can be routed over an APPN network and received by APPN application hosts (end nodes). VTAM is the DLUS. The 3746-9x0 with NNP is the DLUR.

Subset technologies that have been developed to refine the value of APPN include:

- Extended Border Node (EBN)
- Enterprise Extender (EE)
- Branch Extender

Extended Border Node

Extended Border Node (EBN) provides an APPN alternative to SNA Network Interconnect (SNI) that is subarea based. This enables APPN nodes in different APPN networks to communicate. For configurations migrating from large SNA subarea networks to APPN/HPR, where the implementation or design is not complete, there are now new considerations. For large networks, the replacement of SNI with EBN in VTAM is a natural migration path.

Enterprise Extender

Enterprise Extender (EE) is another innovation from IBM to transport SNA traffic over an IP backbone network using routers or IP switches. EE is now an open standard (RFC 2353) and is available from two manufacturers to date (IBM and Cisco Systems). Enterprise Extender uses SNA's HPR technology in which the Rapid Transport Protocol (RTP) ensures reliable delivery and provides efficient flow control between the remote EE endpoint and the application on the S/390. This allows non-disruptive session routing for SNA sessions that spans both the IP and SNA routed networks. Enterprise Extender remote devices may optionally include DLUR and Branch Extender. S/390 Servers with Enterprise Extender can receive HPR traffic over IP transport and thus, may take full advantage of OSA-Express. There are no single points of failure anywhere between Enterprise Extender in the branch and the SNA application if alternative paths exist.

[2] APPN Implementers Workshop

In comparison with DLSw, which uses a TCP connection over the IP backbone, Enterprise Extender sends the SNA traffic into the IP network using the UDP protocol. EE maps the SNA CoS priority bit to a UDP port. The CoS priority bits distinguish between high, medium and low traffic. Most routers can priooritize the UDP port number, and therefore, SNA traffic can be differentiated from other traffic in the network. Note that Enterprise Extender does not require a downstream HPR implementation. Dependent LU requester (DLUR) can be implemented on the Enterprise Extender node, with no changes to the downstream SNA terminal. These benefits bring significant advances in availability and scalability for access to mission-critical applications.

Where the Enterprise Extender has been selected to transport SNA and/or APPN/HPR traffic over an IP network, software solutions for remote network nodes are available from IBM, including Communications Server for Windows NT and Client Server/2.

Branch Extender

The Branch Extender is APPN technology that is designed to run on a remote branch gateway or router. The Branch Extender has the advantage of reducing network topology traffic by presenting an end-node appearance to the network. It can reside on a remote network node or end node and optionally can provide DLUR function. However, when Branch Extenders are cascaded, only the node nearest to the network can perform DLUR.

For configurations migrating from large SNA subarea branch-based networks to APPN/HPR, the Branch Extender provides a solution where the EBN is not available or not appropriate. Software solutions for remote network nodes are available from IBM, including Communications Sever for Windows NT and Client Server/2.

3746 APPN/HPR Summary

The Network Node Processor (NNP) of the 3746-9x0 provides intelligent routing of APPN/HPR and of IP traffic, for connectivity to S/390 servers, and complementing the Parallel Sysplex configuration. The 374x family fully supports SNA subarea, APPN/HPR, and IP environments, thus supporting the S/390 for complete integration into the e-business model.

The 3745 and 3746 are the prime platforms for SNA subarea to APPN/HPR migration in the S/390 server access environment. Recent enhancements to the management of APPN/HPR networks are provided in the 3746 Extended Functions 4, available since July 1999.

Telnet 3270E

The Telnet 3270E (TN3270E) function allows an end user to access SNA applications from a TCP/IP workstation. There are two components of TN3270E:

- The client function that operates on the end user's workstation. The client provides a 3270 session appearance to the end user, but carries the data over a TCP connection.

- The server function that operates somewhere between the TN3270E client and the SNA application on the

  S/390 server in the network. The server function provides an LU for each TCP connection from a remote TN3270E client which, in effect, maps the TCP connections into SNA sessions. The SNA sessions are then carried over an SNA backbone network in the normal way.

  A key consideration when deciding to use TN3270E as a technique to transport SNA sessions over an IP backbone network is where to place the TN3270E server function. There are four possible choices:

- On the S/390 server with the TN3270E server function provided by TCP/IP in OS/390 software. In this case, the SNA-to-TCP conversion is done on the S/390 server. All network traffic across the channel is then carried as TCP/IP.

- On the channel gateway. The SNA-to-TCP conversion is done on the gateway and all network traffic across the channel is carried as SNA. TCP/IP is used from the gateway out to the network and to the remote client.

- At the local site on a LAN-attached server device. The SNA-to-TCP conversion is done locally on this separate TN3270E server device. Network traffic from the TN3270E server is carried across the LAN to the 374x controller, and from there through a channel to the S/390 host.

  Note: This method, unlike the others described above, uses NCP to offload host cycles from the S/390 server. TCP/IP is used from the server out to the network and the remote client.

- At a remote site, which is usually the branch concentrator or a router. SNA sessions, in this configuration, are used from the S/390 out to the remote TN3270E server. TCP/IP is used from the server across the LAN to the TN3270E client.

  Although the most convenient and frequently the most cost-effective deployment is to place the TN3270E server in the same S/390 host as the SNA application. In making the placement decision, you must inventory skills and existing systems, and assess the evolution of existing applications. The placement tradeoffs to consider include price and performance, scalablity, availability and, of course, management tools.

The use of TN3270E as a technique to carry SNA traffic can be enhanced by combining it with Enterprise Extender when the requirement is to have the TN3270E server at the remote branch and an IP backbone network. When combined, the TN3270E and Enterprise Extender together provide a unique solution that not only prioritizes the SNA traffic through the IP backbone, but it also supports HPR out to the TN3270E server. The result is the highest availability attainable for the Parallel Sysplex configuration.

NCP Integrated Product Support

A number of licensed programs are integrated into the load module of NCP to provide support for additional protocols. Although specific to certain environments, these programs are very important for those customers using them, and there tend not to be many replacement alternatives. The main products in this category are:

- *Emulation Program (EP)*: This supports BSC and start-stop protocols. BSC is used mainly for RJE 3780 devices although some networks might also use it for BSC3270. Start-stop is used for older ASCII terminals. The EP can be used when integrated into the NCP load module, in which case the load module is referred to as a *Partitioned Emulation Program or PEP.* Alternatively, the EP might be running stand-alone in a 37xx controller. An alternative to using EP is to use a separate hardware protocol converter.
- *Network Packet Switched Interface (NPSI)*: This provides X.25 support to connect to X.25 networks, or X.25 support on point-to-point links. NPSI provides support for a number of different protocols that have to be considered in two groups:
  - *QLLC*: This provides transport for SNA applications over an X.25 link. NPSI is supported on the 3745 and 3746 today. Support for QLLC is also provided by a number of alternative products. For example, the 3746 has a licensed internal code solution for X.25 QLLC support (the X.25 support feature), which removes the requirement for NPSI in the NCP.
  - *Non-SNA protocols*: NPSI supports a number of other protocols including PCNE, GATE, DATE, X.3 PAD, and so on. These tend to be used for very specific applications, maybe involving custom-built terminals, and they are very commonly used by NPSI customers. There is no easy migration option for many of these protocols without rewriting the application or choosing different hardware platforms for the terminals. Therefore, the recommendation for X.25 non-SNA protocols is to continue to use NPSI.
- *Network Terminal Option (NTO)*: This provides support for ASCII or asynch connections. Like the EP, NTO cannot be replaced by a router, and needs an external protocol converter to replace its functionality.
- *Other programs* such as Network Routing Facility (NRF), Network SNA Interconnection (NSI), MERVA , and so on also have no direct replacement in hardware or software.

3745/3746 Evolution to e-business Scenarios

There are many environments where the 3745 and 3746 controllers are used today, and it is impractical to list all the possibilities. Therefore, the following discussion is limited to the most common configurations and associated 374x scenarios.

The majority of NCP placements are in three distinct environments:

1. Classic SNA subarea with or without NCP Extended Options (X.25 NPSI, EP, NRF, and so on)
2. Classic SNA with emerging APPN
3. SNA/APPN converging with IP networks

Today, we will cover ten scenarios, each having a different configuration and possible technology options. Scenarios1 through 5 involve classic SNA while scenarios 6 through 10 involve converging SNA and IP networks. In summary, these scenarios are:

Scenario 1: SNA subarea and NCP Extended Options (NEO) products

Scenario 2: SNA subarea with multiple 3705/3725/3745s (data-center consolidation)

Scenario 3: Replacement of remote 37xx Communication Controllers

Scenario 4: Large SNA/APPN/HPR network

Scenario 5: SNA subarea-to-APPN/HPR migration

Scenario 6: Parallel Sysplex with subarea SNA network

Scenario 7: Parallel Sysplex with subarea network and TCP/IP

Scenario 8: SNA Subarea and separate TCP/IP network

Scenario 9: Parallel Sysplex/network consolidation

Scenario 10: Access to Transaction Processing Facility (TPF)

The characteristics of the ten scenarios and the recommended actions are summarized in the table.

Characteristics of the Ten Scenarios and the Recommended Actions

Scenarios	SNA Subarea	APPN/ HPR	TCP/IP	Parallel Sysplex	Planned Direction
1	yes EP, SNI				Retain the 3745/3746 and NCP to support other NCP-related programs.
2	yes 37xx				Consolidate older controllers for leased line concentration onto the 3746.
3	yes Remote 37xx				Replace with FRADs where appropriate and with replace ment 3745-170 if this is not possible.
4	yes	yes			Upgrade the 3746 with Extended Functions 4 to enhance perfor mance without requiring hardware upgrades.
5	yes SNI	yes			Migrate to EBN where appropriate, and continue with SNI if EBN is not possible.
6	yes				APPN/HPR migration in network
7	yes	yes	yes	yes	IP backbone using the 3746 for SNA and OSA-Express for TCP/IP.
8	yes		yes		IP backbone using 3746 and DLSw downstream for SNA and separate router for TCP/IP.
9	yes	yes	yes	yes	IP backbone using Enterprise Extender for SNA and OSA-Express for all connectivity.
10	yes TPF		yes TPF		Migrate from Offload and use the 3746 for both SNA and TCP/IP.

Classic SNA

Scenario 1:
SNA subarea and NCP Extended Options (NEO) products

Current Environment

- A number of networks are supported, including a SNA backbone network, an EP network, and an X.25 network.
- EP is used to handle older protocols, including BSC3270.
- NPSI is supporting a combination of SNA (QLLC) and non-SNA (non-QLLC) protocols.

Current Plans

- In an attempt to cut costs, alternatives for the 3745/3746 and NCP are being considered.
- The EP applications are being rewritten using SNA.
- Options for the BSC3270 terminals are being considered.
- Alternative plans are being considered for the NPSI non-QLLC protocols.

Considerations

- EP has no direct replacement in the router environment. Although some external protocol-converter solutions are available, it makes little sense to do this, given the investment in the 3745/3746. Another option is to rewrite the applications on the S/390. In this scenario, it is assumed that the applications will eventually be rewritten. Therefore, it makes no sense to consider hardware solutions at this stage. Waiting for the new applications will be the most cost-effective option in the long term.
- The BSC3270 terminals can be handled by a frame-relay access device (FRAD) if that is appropriate. However, an alternative solution, once the S/390 EP applications are rewritten, is to use the NCP's BSC3270 support.
- X.25 NPSI QLLC in the 3745 could be replaced by the X.25 support feature in the 3746 processors.
- NPSI non-QLLC is very hard to replace without purchasing different terminals, or rewriting applications. It probably makes sense to keep the 3745/3746 for NPSI use until the applications have been rewritten or the network has been changed.
- Other additional NCP programs and protocols may also be in use:
  - The Airline Line Control (ALC) protocols are fully supported by the 3745. There are no alternative solutions from IBM, although Cisco routers do offer an ALC solution. However, given that most of the other options require the continued use of the NCP, it is not recommended to install a router purely for this purpose.
  - For all other NCP additional program products (NRF, NSI, NTO, and so on), the best solution is to continue to use the products (see the discussion in "Technology Overview"), because very few general-purpose solutions exist.

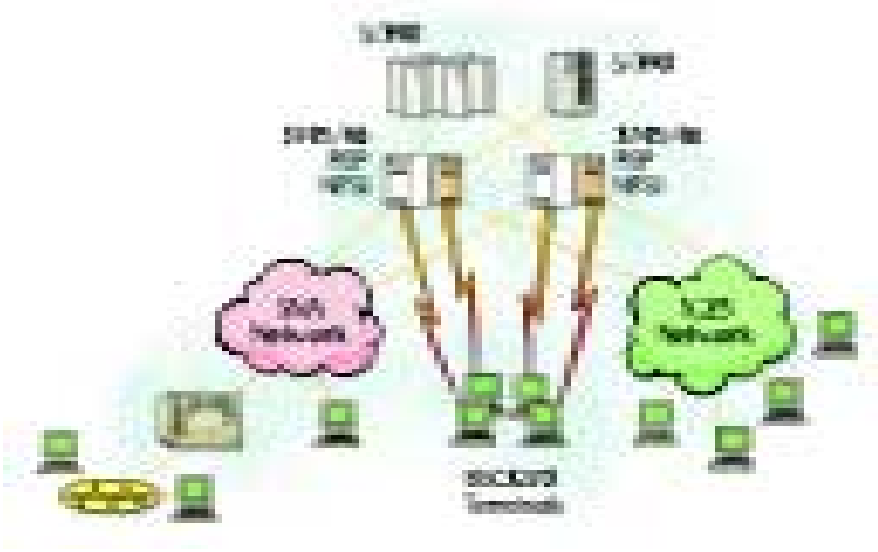

Recommendations

- Keep the 3745/3746 and NCP for SNA subarea traffic.
- As the S/390 applications are rewritten or replaced, EP can be removed from the NCP load module.
- Move the BSC3270 terminals to be controlled by the NCP.
- Because there is a requirement for both NPSI QLLC and non-QLLC, the best solution is to continue using NPSI. In addition, moving the NPSI QLLC to the outboard adapters on the 3746 may enable you to reduce NCP Tier Level and reduce both NPSI and NCP license charges.
- When a solution is found for the NPSI non-QLLC applications, then the QLLC traffic can be moved to use the 3746 support, and NPSI can be removed from the NCP load module.
- For all other NCP additional program products, the best solution is to continue to use the NCP Extended Options (NEO).
- Retain the 3745 for any ALC Support.

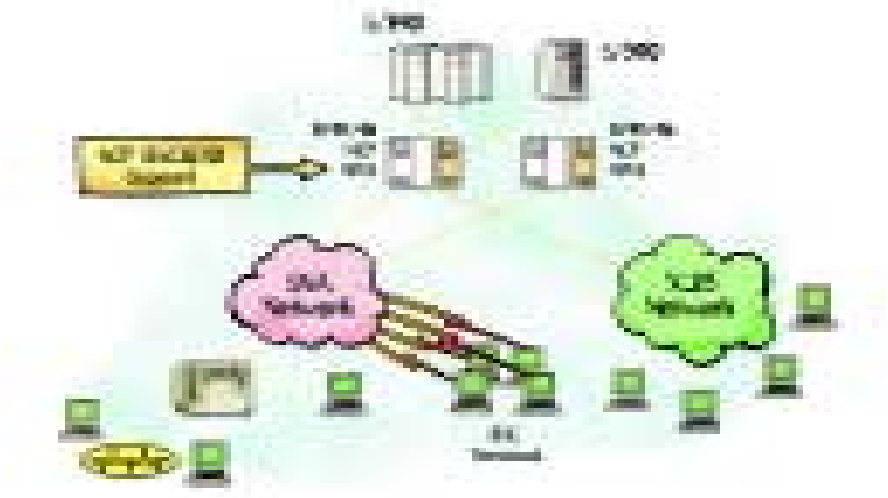

Scenario 2:
SNA subarea with multiple 3705/3725/3745s (data-center consolidation)

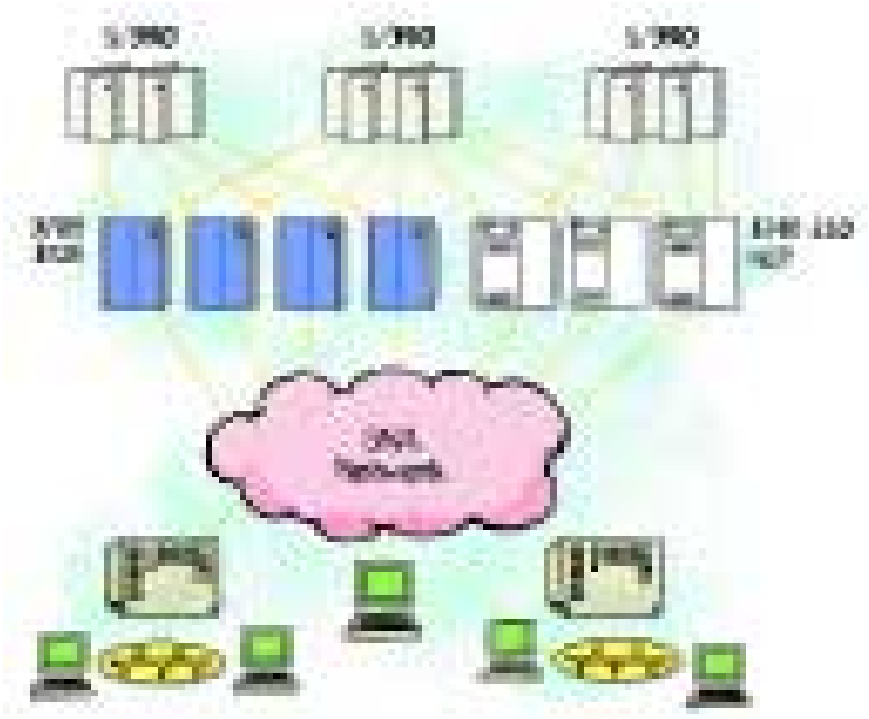

Current Environment

- The current network is based on subarea SNA, using four 3705/3725 communication controllers at the data center for historical reasons.
- In addition, there are three 3745-210 controllers.
- The seven controllers are handling large numbers of serial lines, and there are no plans to move to an IP infrastructure.

Current Plans

Maintenance on the 3705 and 3725 controllers has already been withdrawn. (The assumption is that this configuration is in the U.S., where maintenance for the 3705, 3720 and 3725 was withdrawn on January 31, 1999.) An alternative platform must be found.

Considerations

- The older controllers are running PEP (a combination of NCP and EP in the same load module) and no alternative hardware solution has been found. The old BTAM applications are being rewritten.
- The only viable solution is to provide EP access to the S/390 BTAM applications using another PEP solution.
- Consolidating the lines from the four older controllers will require additional ports and it is recommended that two 3746s be installed. Using the 3746 for consolidation and NCP program Tier C (there are no NCP charges for the 3746) will be a more cost-effective solution. Any new or existing 3746 is eligible for the ordering of a new Tier C license, provided that an NNP Type 3 (#5053) is configured with the 3746.

Recommendations

- The three 3745-210s must first be upgraded to 3745-31As.
- The second step will be to install the two new 3746s and move the 3745 connections to the 3746 adapters.
- Finally, all the network attachments from the 3705 and 3725 controllers can be moved across to the 3746.

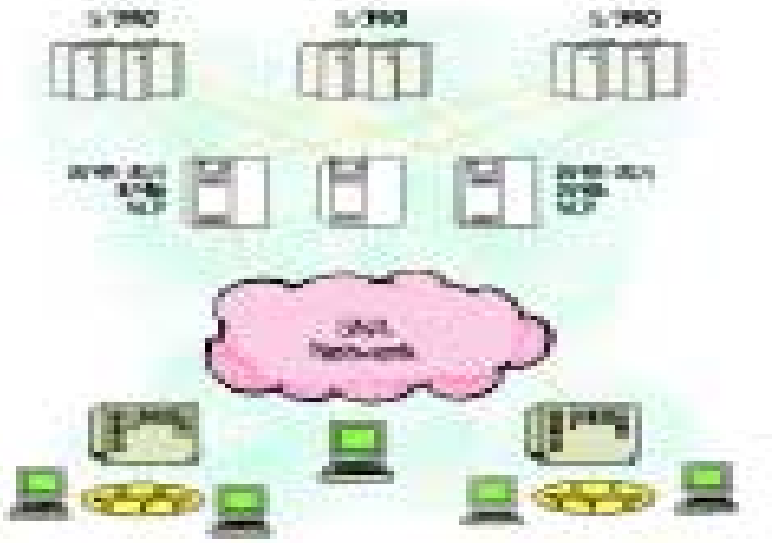

Scenario 3:

Replacement of remote 37xx Communication Controllers

Current Environment

- The installed network is based on SNA subarea.
- The network comprises six remote 3720 Communication Controllers at the six regional centers. Each 3720 has two upstream SDLC links, one to each 3745 in the data center. The remote FEPs are used for SLDC and token-ring concentration.

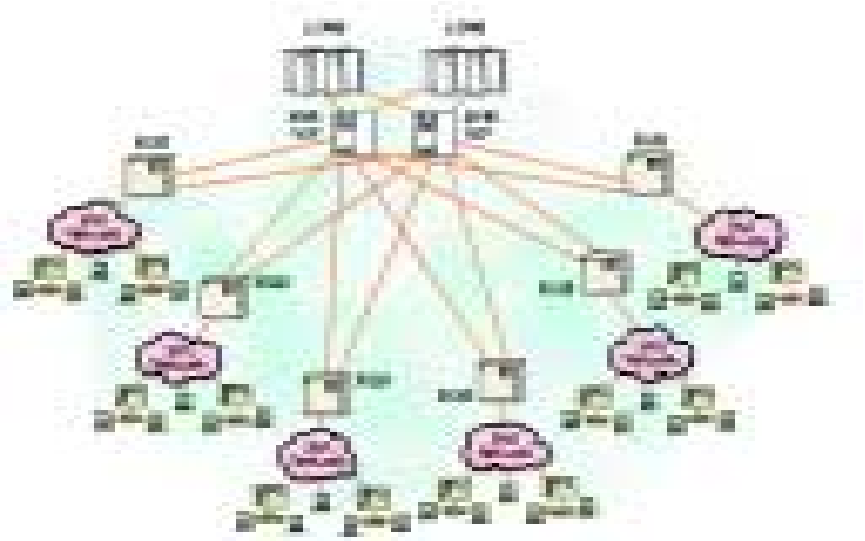

Current Plans

- Maintenance on the 3720s will soon be withdrawn. (The assumption is that this configuration is in EMEA where maintenance will be withdrawn on January 31, 2001.)
- An alternative solution must be implemented in 2000, before the normal year-end freeze, which starts in October 2000.

Considerations

SDLC and token-ring consolidation can be provided by a number of alternative solutions, including:

- Replace the 3720s with 3745-170s on a one-for-one basis. This will be the simplest solution, but might be expensive in terms of hardware and NCP costs. However, it does also provide a solution in the case where some of the NEO products are used (such as NPSI or NSI). In these cases, this may be the only practical option.
- Replace the 3720s with frame-relay BAN devices, such as FRADs. The upstream link from the region will be frame relay on a point-to-point link. The local data center NCPs must be at NCP V7R3 or later. No VTAM system changes are required. Replace the 3720s with APPN routers, using DLUR to support the 3270 traffic. This will require an APPN migration, which may not be appropriate in all cases because it involves VTAM changes.

Recommendations

- Replace remote 3720s with frame-relay BAN devices where possible, and implement frame-relay links on the local data-center NCPs.
- In other regional centers where NCP cannot be removed, consider the acquisition of a 3745-170 or 3745-17A.

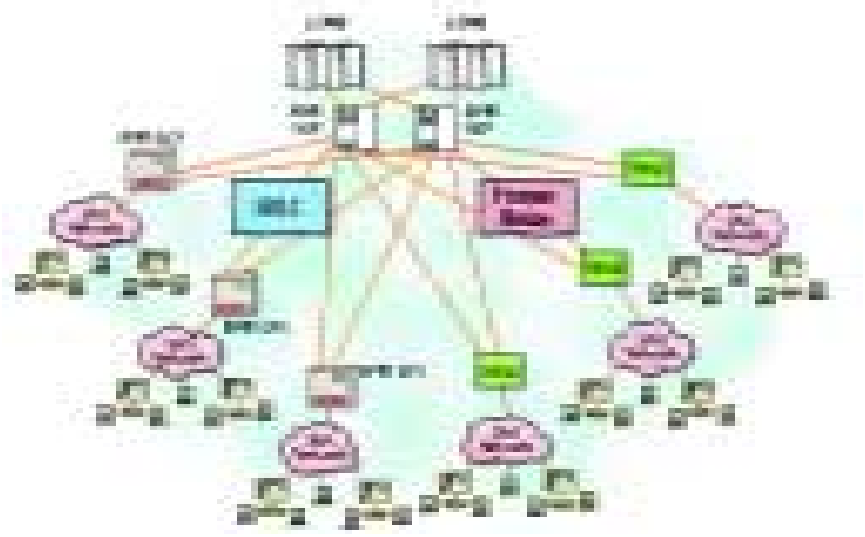

Classic SNA Evolving to APPN/HPR Networks

Scenario 4:
Large SNA/APPN/HPR network

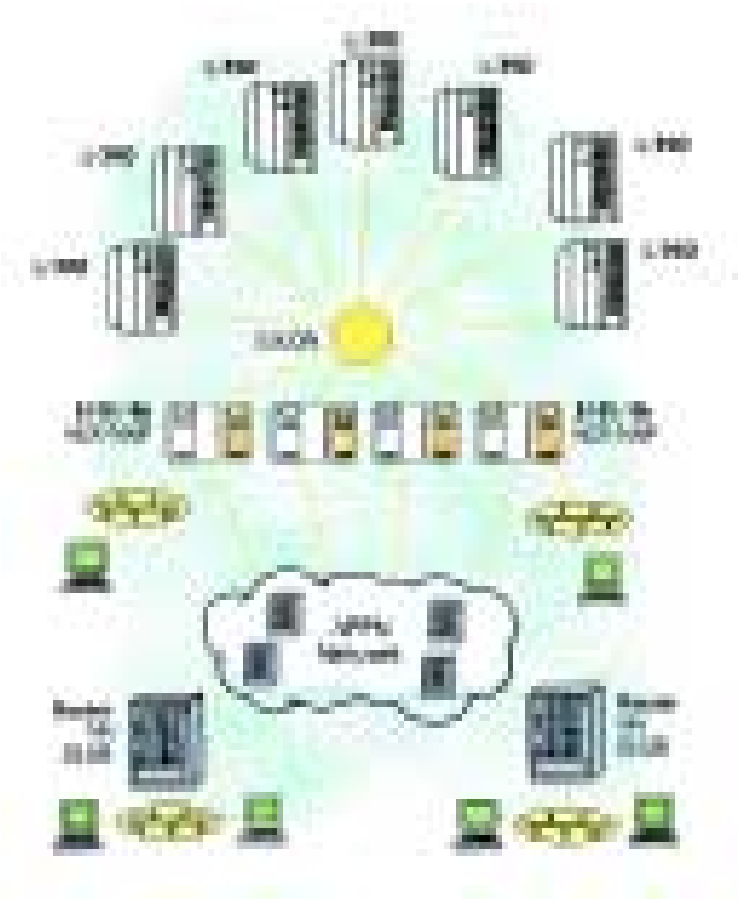

Current Environment

- A large SNA network is accessing the S/390 server. The 3745 and 3746 are installed at the data center.
- The 3746 is controlled by NCP and NNP. NCP controls the subarea SNA traffic, which is located in the data-center location. The Network Node Processor (NNP) controls the Advanced Peer-to-Peer Networking/High-Performance Routing (APPN/HPR) traffic that comes from the remote network.
- The remote network is connected via token-ring ports to the backbone router infrastructure.
- All remote workstations are router connected. Performance, response time, and availability are key factors for remote users. DLUR is used on the remote routers for subarea access.

Current Plans

- Optimize network performance and capacity.
- APPN/HPR is being implemented in certain remote sites to support applications with high-availability requirements.
- All SNA traffic is growing in terms of volumes and number of users.

Considerations

- The latest version of the 3746 microcode, delivered in July 1999, provides token-ring connection balancing for SNA and APPN/HPR users. This code is delivered as part of 3746 Extended Functions 4. It allows connections to be balanced over multiple token-ring ports and across multiple 3746s.
- This microcode feature can be installed without any change to the installed Token Ring Type 3 processors. The new function results in a more controlled distribution of connection requests by having duplicate TICs delay their respones to route-discovery requests.
- SNA traffic is growing and the 3745 central control unit (CCU) is heavily loaded. One concern is whether it can handle the expected growth. However, as the APPN/HPR migration continues, certain sites are moved away from the NCP and connected to the 3746 NNP, thereby saving the corresponding 3745 CCU load. This means that the 3745 will be better placed to support the SNA subarea growth.
- The 3746 NNP is performing high-throughput APPN/HPR ANR routing, and is therefore well placed to handle additional traffic volumes. In addition, the scalablility of the 3746 means that additional processors can be added to support more connectivity as the number of users increases.
- If the network growth is likely to require Fast Ethernet, ATM, or ISDN connectivity in the future, then the Multiaccess Enclosure (MAE) should be considered for the 3746.

Recommendations

- Upgrade the installed 3746 controllers with 3746 Extended Functions 4 to provide token-ring connection balancing.
- Continue the rollout of APPN/HPR, optimizing performance and availability for SNA applications, and using the 3746 to handle traffic growth.
- Use the MAE to support high-speed connectivity options.

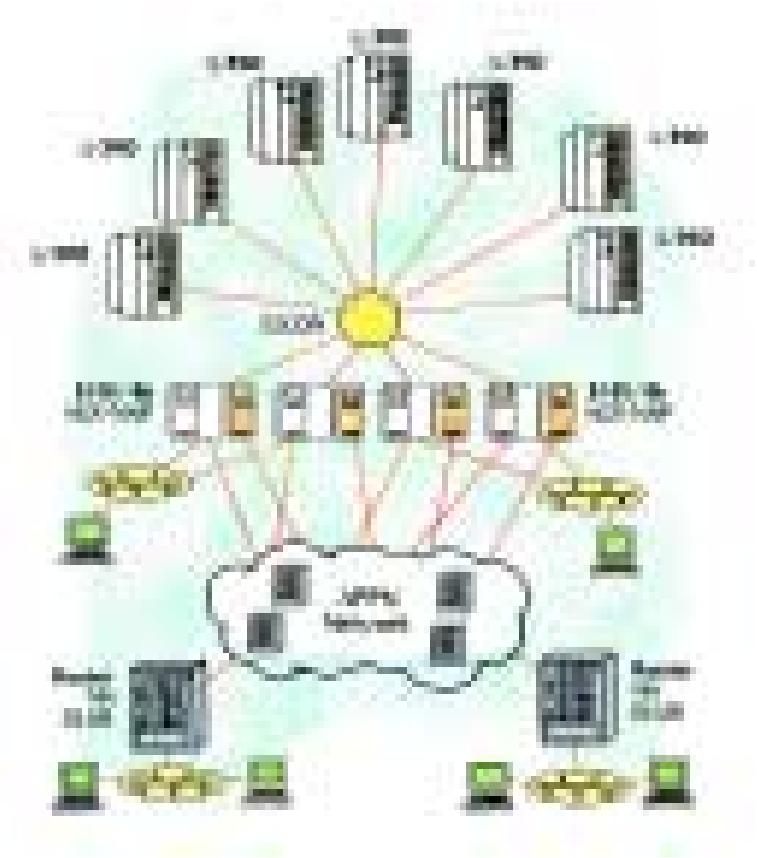

Scenario 5:
SNA Subarea to APPN/HPR migration

Current Environment

- SNA network supports a combination of SNA subarea and APPN/HPR traffic.
- Network gateway to the S/390 server is 3745/3746, with NNP to support APPN/HPR traffic.
- The network is implemented using router infrastructure and frame-relay protocol.

Current Plans

- SNA subarea to APPN/HPR migration continues, with the objective of migrating away from NCP completely.
- SNI is currently used by various partners, and Extended Border Node (EBN) is being considered as an alternative.

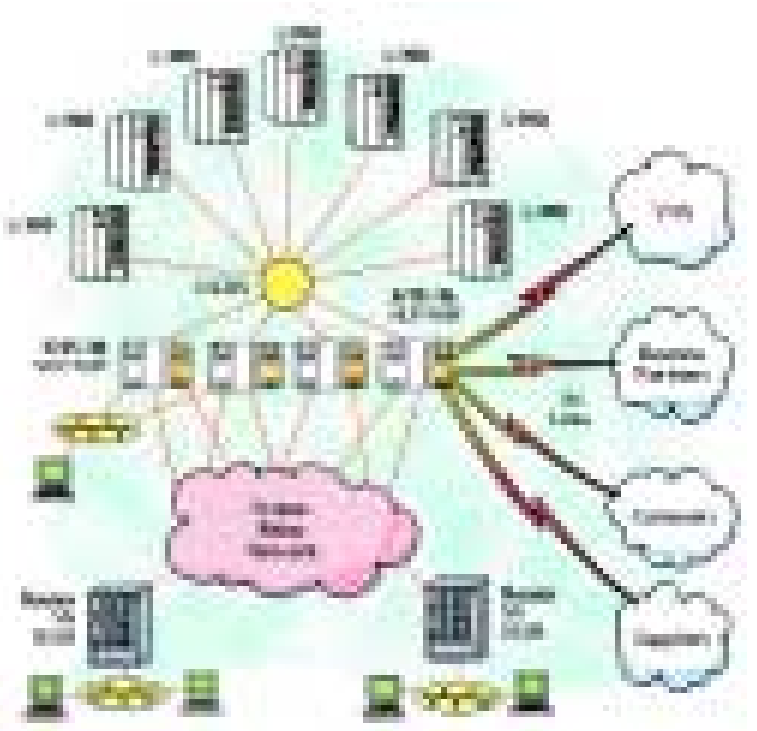

Considerations

- Cisco has no EBN solution. Therefore, VTAM/NCP is the only solution.

- If the configuration requires an APPN border inside the network, the VTAM/NCP EBN solution is not practical. One of the other IBM implementations (such as Branch Extender) should be used for this requirement.

- Most SNI users today are communicating S/390-to-S/390, and so the VTAM EBN is a very practical option. One of the characteristics of EBN is that both ends of the link must be using APPN. In the mid-1990s, when EBN was first shipped with VTAM, implementation was not always possible because many companies were not at the correct levels of VTAM to support APPN. Today, most organizations that use SNI are at the latest levels of VTAM for Year 2000 reasons, and are APPN-capable, if not actually using APPN.

- The implementation of the VTAM EBN does not mean that the whole network must be migrated to APPN, which has been a concern in the past. The VTAM APPN implementation provides an interface between subarea networking and APPN. This makes it possible to retain a subarea SNA network, with only the SNI links migrated to EBN links. This may be desirable for the partner networks, who may not be able to make any network changes.

- An alternative to the EBN is the Branch Extender, which resides in a branch. It does not lend itself to replacing EBN inside the network, because it cannot act as an intermediate APPN/HPR routing node. Most SNI configurations are for S/390-to-S/390 links, to connect to business partners, suppliers, or value-added network (VAN) providers. The Branch Extender will not support these S/390 configurations. It is best placed to handle the case where a company network needs to be split into separate subnetworks to reduce topology updates in the network, while maintaining full APPN directory services. So, Branch Extender is not suitable in this scenario.

- Another alternative for the replacement of SNI is to persuade the partner to take over responsibility for the SNA gateway. Most SNI links today are configured to use the back-to-back gateway configuration. This configuration provides the maximum amount of isolation and control and allows for the maximum number of network addresses. However, there is another way of defining the SNI link, using the single gateway configuration. This might involve a considerable amount of reconfiguration and testing, but may be appropriate for some environments. The steps required are:
  - Persuade the SNI partner to change their back-to-back definitions to single gateway definitions.
  - Configure an SNA passthrough solution at this end of the SNI link. This can be implemented using OSA-Express. Alternatively, a router can be used. If the SNI link today uses a serial line, then this line must be terminated by a device such as a router because the OSA adapter does not support serial lines.
  - Remove all SNI definitions from NCP and move the connection to the OSA-Express or router.

- In this scenario, the SNI links are company-to-company (S/390-to-S/390). It is recommended that EBN be used wherever possible. However, it is assumed that not every partner will be able or willing to migrate to APPN. So, a combination of EBN and SNI is recommended.

Recommendations

- Complete the APPN migration in the network.
- If EBN migration has already started, continue with current product rollout, including 221x for EBN.
- In this scenario, because the rollout of EBN has not yet started, use VTAM EBN for S/390-to-S/390 traffic where possible.
- Where SNI partners may not be able or willing to migrate to APPN concurrently, these links must stay as SNI, which requires a gateway NCP.

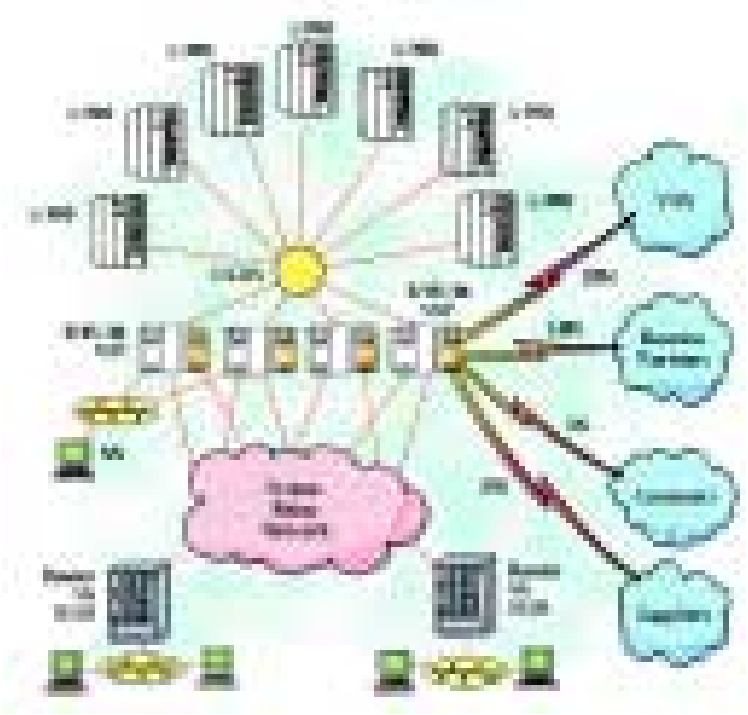

Converging SNA and IP Networks

Scenario 6:
Parallel Sysplex with subarea SNA network

Current Environment

- An SNA subarea network is carrying 3270 application traffic into the S/390 Parallel Sysplex.
- The 3745/3746 is under NCP control today.
- No TCP/IP is installed or planned as part of this network.

Note: WLM is OS/390 WorkLoad Manager, which provides resource-utilization information in the Parallel Sysplex configuration.

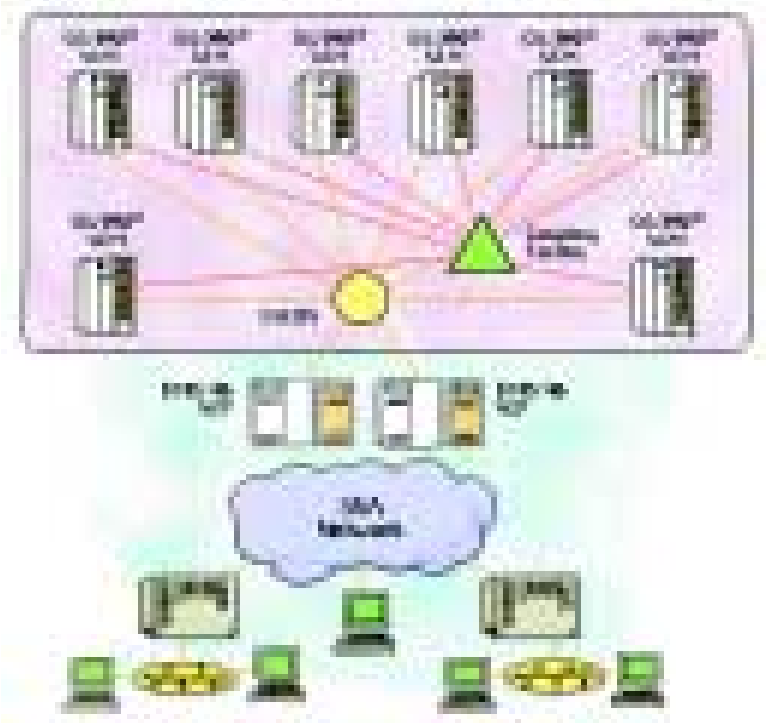

Current Plans

- APPN/HPR is required end-to-end, to provide high availability across the network for all 3270 traffic.
- IBM routers are currently selected for APPN NN and DLUR.
- Upgrade 3746 with Network Node Processor (NNP) to provide APPN NN gateway to S/390.

Considerations

- The 3746 NNP will be used as an ANR router for HPR traffic. This HPR traffic is carried from remote network nodes acting as RTP endpoints, over the ESCON channel to RTP endpoints in VTAM.
- APPN/HPR transport will be provided natively by the router, using frame relay, X.25 or SDLC.
- If the rollout plans change and remote routers are used without RTP and DLUR functions, the 3746 can be used to fulfil this role instead of acting as an ANR router. The SNA traffic will be bridged across the router network, and the 3746 NNP will act as network node, with DLUR and HPR (RTP endpoint) functions.
- In the Parallel Sysplex environment, the latest technology is used in the S/390 server, and so it is very likely that OSA will be available, if not today, then within the medium-to-long term. In this environment of no IP on the S/390 and a growing SNA (HPR) network, OSA-2-based token-ring links might be used to add additional link capacity for HPR and subarea traffic between the 374x and the S/390 server. However, the NCP function for subarea traffic and the ANR, RTP and DLUR functions for HPR traffic remain essential components of the SNA network solution.
- One consideration could be to replace the 374x with a channel-attached router. This might be attractive for the HPR environment if the router:
  - Has robust ANR, RTP and DLUR function
  - Has native HPR transport over ESCON
  - Has the session capacities essential to support your environment

  Note that the 3746-9x0 with NNP routes HPR, APPN and IP, and thus is a router.

Recommendations

Note: Remote routers can be 2210s, 2212s, or 2216s.

- Make no changes to the data-center strategy. Keep the 3745/3746/NCP for subarea traffic and upgrade the 3746 with the NNP for DLUR. Migrate the SNA 3270 traffic from NCP control to NNP control with DLUR as the network migrates to APPN. The 3746 is highly scalable and well positioned to accommodate future network growth. Retaining the 374x products ensures smooth migration and continued high availability to your mission-critical applications.
- APPN network designs should consider the incorporation of the Branch Extender at remote sites to reduce the number of network nodes. The resulting benefits include improved network efficiency with smaller APPN topologies and dramatic increases in network scalability.
- Continue with current plans for network rollout. Due to the end of marketing of the 221x family anticipated for 4Q 2000, the 221x units required to complete the deployment plans should be acquired in 2000.

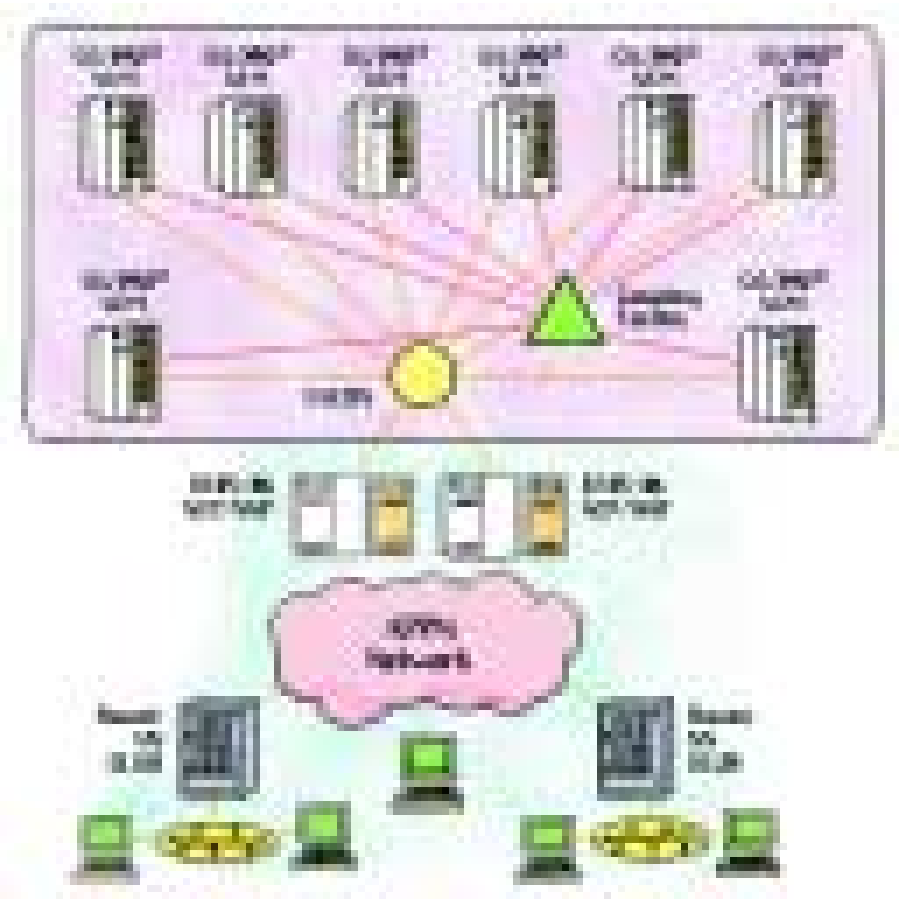

Scenario 7:
Parallel Sysplex with subarea network and TCP/IP

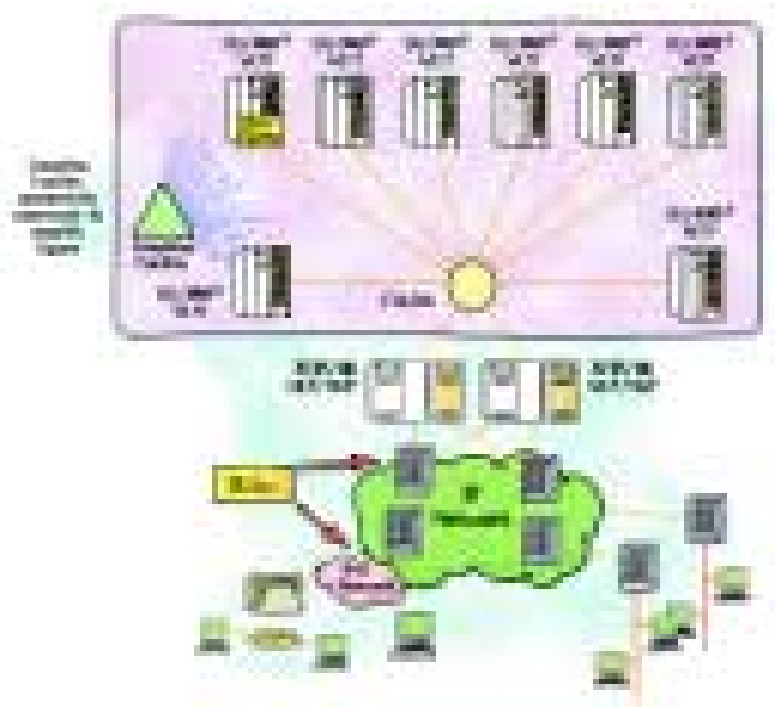

Current Environment

- The SNA subarea network is carrying 3270 application traffic into a S/390 Parallel Sysplex.
- The 3746 is under the control of both NCP and the NNP.
- Some APPN is used inside the Parallel Sysplex for generic resources (GR), providing workload balancing. Users of GR applications are mainly attached to the Parallel Sysplex through the NCP. The APPN network node is provided by the combination of VTAM and NCP, as a composite network node.
- Some APPN/HPR is used inside the Parallel Sysplex for multinode persistence sessions (MNPSs), providing high availability. Users of the MNPS applications are controlled by the 3746 NNP. These users also have access to GR applications.
- Significant and fast-growing TCP/IP traffic carried into the Parallel Sysplex through the 3746. S/390 used as main Web-server platform for Internet and intranet applications.
- The SNA traffic is carried over the IP backbone network using DLSw.

Current Plans

- Continue to grow the APPN/HPR network to enable efficient access to Parallel Sysplex SNA applications, which are shifting toward native APP and HPR. In the 3746-900, shift workloads from NCP to NNP with DLUR enabling NCP consolidations.
- The Parallel Sysplex is expanding to handle the increased traffic load, especially for the Web server applications, and two new S/390 servers are to be implemented.
- OSA-Express will be implemented during 2000 for TCP/IP traffic only. This will add significant IP link capacity to the S/390. Connectivity will be through Gigabit Ethernet.
- As TCP/IP traffic volumes increase, Differentiated Services on the OS/390 will enable QoS across the IP backbone.

Considerations

- The 3746 as APPN/HPR gateway to the S/390 provides essential function (RTP, ANR, DLUR, and so on) and value to the APPN/HPR network, and remains a critical platform in the network. The 3746-900 with NNP provides the smoothest migration from subarea (NCP-controlled) traffic to APPN /HPR (NNP-controlled) traffic, which is fundamental to maintaining high availability of the mission-critical SNA applications during the migration period.
- The SNA 3270 traffic is carried over the router network using DLSw. The router terminates the DLSw connection and passes the SNA 3270 traffic over token ring into the 3746. The NNP converts the subarea SNA 3270 traffic to APPN using DLUR, and provides the RTP endpoint for the HPR connection to the correct S/390 server, primarily over ESCON channels. With the availability of OSA ports that support native SNA traffic, connectivity between 374x and S/390 might migrate from ESCON to LAN links. VTAM provides the RTP endpoint function on the S/390 server.
- For access to 3270 applications on the S/390 there is a clear shift on the desktop (client) to TN3270E. The rate and pace of this trend will vary depending on the reliability of the aging desktops and with the demand for access to Web applications from these desktops. However, the direction is undeniable. IT designers must establish a deployment strategy for TN3270E servers or gateways.

Based on current characteristics, including scalablity, capacity, availability, and price, the clear choice has emerged only recently as the S/390 with OS/390 V2R8 or later. The latest version of OS/390 with its many enhancements to the TCP/IP stack makes the TN3270E server function on the S/390 an attractive proposition. The ultimate solution for TN3270E access will be to bring the IP traffic into the S/390 over the OSA-Express. As a migration step toward OSA-Express deployment, the IP routing function of 3746 and the ESCON links can be used to provide access to the TN3270E Server on OS/390.

- For access to IP applications on the S/390, the performance of OSA-Express with Gigabit Ethernet will be superior to ESCON. The OSA family, with OSA-Express being the latest and greatest member of the family, will be the strategic network access technology for the S/390 servers. As OSA-Express deployment expands from the first shipments in 1999, OSA-Express will emerge as the preferred network access for IP. This will be due to the dramatic gains in link capacities (Gigabit Ethernet vs. ESCON) and in price-per-MB.

  The latest TCP/IP functional enhancements also provide advantages for OSA-Express compared with ESCON. For S/390 servers that are not OSA-Express capable, 3746 provides a proven platform with ESCON links. 3746-9x0 with NNP supports RIP V1 and V2 and Open Shortest Path First (OSPF) routing and is optimized to interoperate with the Routing Information Protocol (RIP) and OSPF capabilities of OS/390.
- For subarea SNA traffic, there is no compelling reason to replace the 3745/3746/NCP. In fact, there are clear disadvantages in doing this, including:
- An increase in S/390 MIPS due to moving the SNA boundary function from NCP to VTAM.
- Indirect routing—all the 3270 traffic will pass through an additional S/390 VTAM hop (where the boundary function is located).
- Degradation of the network management and operational capabilities.
- Loss of recovery and alternative paths to SNA applications using techniques such as token-ring connection balancing.

  Note: OSA-2 with a token-ring attachment to the 374x, replacing ESCON links, can be an alternative solution without realizing the above-mentioned disadvantages.

Recommendations

SNA traffic:

- Continue to use the 3745/3746 for SNA subarea and APPN/HPR access to the S/390.
- Use the 3746 under control of both NCP and NNP.

TCP traffic:

- OSA-Express is strategic for IP network access to the S/390. For OSA-Express-capable S/390 servers, migrate to OSA-Express connected to the IP via a LAN switch.
- For S/390 servers that are not OSA-Express capable, the 3746 IP routing capabilities over ESCON links provide seamless platform integration for both IP and SNA S/390 access.

Current Plans

- There is a requirement to consolidate SNA and TCP/IP traffic onto a single router-based transport network.
- There are no plans to migrate to APPN/HPR.

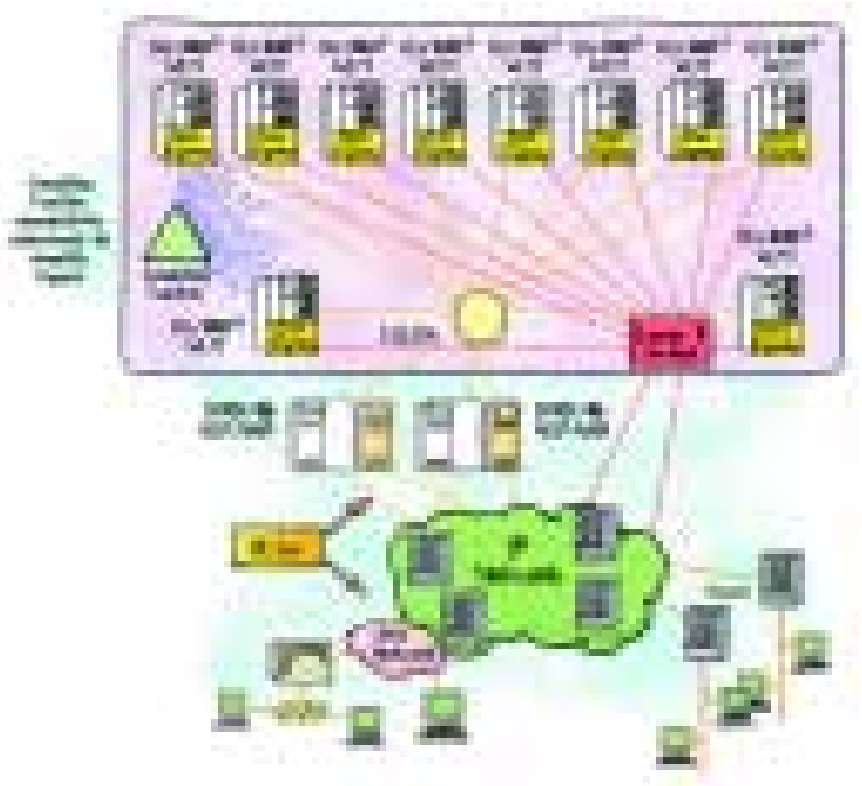

Scenario 8:

SNA subarea and separate TCP/IP network

Current Environment

- An SNA network supports legacy 3270 traffic and there is movement to TN3270E. However, due to the economics of replacing terminals, this movement is expected to be both slow and continuous. SNA 3270 will exist for many years.
- The 3745/3746s are under the control of the NCP.
- The SNA traffic is growing because of new users and business growth.
- A separate IP backbone network carries TCP and LAN traffic.

Current Plans

- There is a requirement to consolidate SNA and TCP/IP traffic onto a single router-based transport network.
- There are no plans to migrate to APPN/HPR.

Considerations

There are a number of possibilities for carrying the SNA 3270 traffic over a router-based backbone network:

- Enterprise Extender (EE) can be used to carry SNA traffic over IP in the router-based backbone. It should be noted that EE uses HPR endpoints for which there are no stated migration. The incorporation of HPR migration should be driven by the data-center organization and their plans to use advanced VTAM function, and not driven by the networking organization.
- TN3270E is an excellent alternative for this environment. TN3270E servers can be configured with SNA subarea or APPN to match your data-center requirements. However, TN3270E requires migration of the 3270 client, which often is staged due to the financial investments required to upgrade or replace thousands of desktops. TN3270E is positioned as a premier technology for 3270 migration to IP. However, deployment of TN3270E clients will take time. See Scenarios 7 and 10 for additional perspectives on TN3270E.
- Data link switching (DLSw) is among the most popular technologies for transporting SNA traffic using IP in the router backbone. Because DLSw is confined to the router platforms, the rollout is managed by the networking organization. From an availability and efficiency perspective, DLSw has limitations. DLSw requires a process-intensive conversion to terminate TCP connections and to initiate SNA link-level protocols for thousands of sessions. Deployment of DLSw in large networks may require hundreds of routers for the termination of DLSw connections. The expense associated with this should be kept in mind.
- In addition, availability issues come into play with the outage of a DLSw gateway that is servicing several hundred sessions. The routers that terminate the DLSw connections are single points of failure. Any outage is disruptive as the sessions are broken, and the reconnect delays might be extended so as to further impair productivity. These issues do not exist with EE. DLSw is implemented in the remote branch router and in a data-center router. The data-center router (DLSw server) terminates the TCP connection and passes SNA traffic to a 374x or directly to VTAM. DLSw is readily available on IBM and other manufacturers' routers.
- Boundary access node (BAN) provides for the sharing of frame-relay links by SNA and TCP/IP natively. Bridging of the SNA traffic between LAN links and frame relay is at the core of BAN. With BAN, SNA can share frame-relay links with TCP/IP. This may be deployed with SNA and IP being segregated to separate permanent virtual circuits (PVCs), each with its committed information rate (CIR). Or SNA and IP can share a PVC and a CIR, with SNA and IP traffic differentiated based on RFC1490 headers. BAN is an alternative to DLSw. BAN is also supported in ATM networks.
- Sharing a data transport such as frame relay, X.25, or ATM may not be possible in all countries, depending on the services available. Alternatively, leased lines operated with frame-relay data-link control can be used. Frame-relay access devices (FRADs) are a class of devices that offer the benefits of native transport without encapsulation.

For this scenario, both DLSw and BAN should be considered. DLSw supports an IP transport network, but might be expensive to terminate large numbers of connections. BAN is supported on a multiprotocol network over a single WAN Layer 2 infrastructure. Either choice is dependent on the router platform for DLSw and for BAN.

As in previous scenarios, there is also no advantage in keeping the SNA subarea infrastructure and simply replacing the 3745/3746/NCP with a channel-attached router. In fact, there are clear disadvantages in doing this, including:

- An increase in S/390 MIPS due to moving the SNA boundary function from NCP to VTAM.
- Indirect routing—all the 3270 traffic will pass through an additional S/390 VTAM hop (where the boundary function is located).
- Unnecessary cost, because it is likely that OSA-Express will become the preferred connectivity option in the medium-to-long term.
  As SNA expansion occurs, additional capacity can be added to the 3746. The 3746 is a extremely scalable platform and can be upgraded to provide additional processing power and connectivity, providing the largest connectivity of any channel-attached gateway.

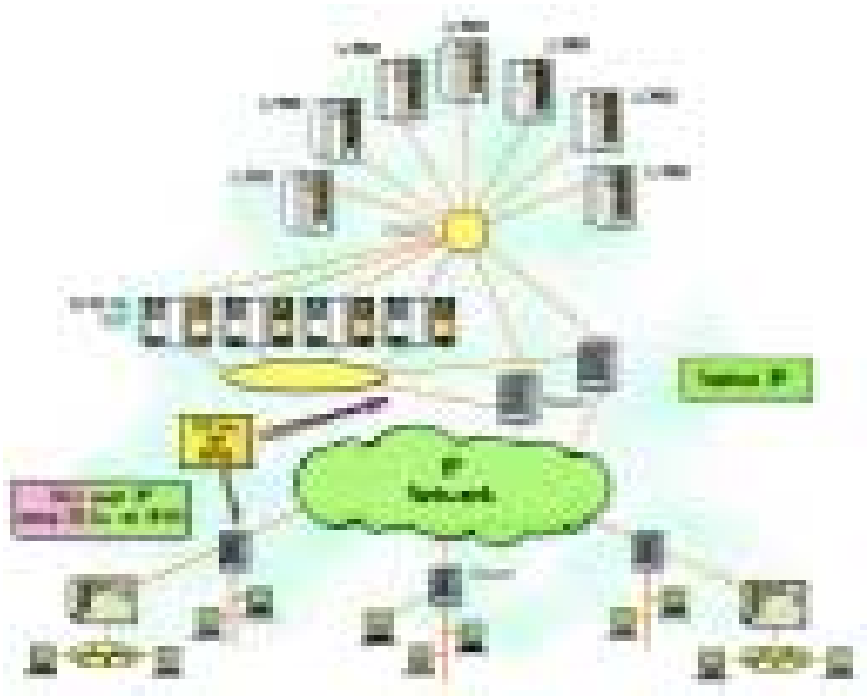

Recommendations

SNA Traffic:

- Continue with existing 3745/3746 for SNA subarea access to S/390.
- Upgrade the 3745 to maximum storage (16 MB) to provide boundary function and address the growing session requirements. This can also facilitate 374x consolidation.
- Expand the existing 3746 when additional connectivity is required (ESCON or token ring for OSA-2 links). Also consider upgrading from TIC-2 to TIC-3 for growth in LU capacity if your 3746 is near capacity.
- Implement DLSw or BAN in routers downstream of the 3745/3746 to bring in the SNA traffic from the router backbone network.

TCP Traffic:

- As the TCP traffic grows, you can add capacity for IP access to the S/390 server with additional routers.
- As OSA-capable S/390 servers are deployed, plan to take advantage of the OSA capacity and functionality and directly connect the S/390 servers to the switched campus infrastructure.

Scenario 9:
Parallel Sysplex/network consolidation

Current Environment

An SNA subarea network is carrying 3270 application traffic into the S/390 Parallel Sysplex.

- Some APPN is implemented inside the Parallel Sysplex for generic resources.
- A 3745 is under NCP control supporting SNA subarea and APPN network node (using VTAM composite network node).
- Separate routers provide TCP/IP traffic access to the Parallel Sysplex. New Web server-based applications require workload balancing.
- The IP backbone network is based on any manufacturers' routers.

Current Plans

- OSA-Express will be implemented for IP during 2000 and connectivity will be through Gigabit Ethernet.
- There are plans to migrate the existing SNA subarea network to the IP backbone concurrent with the migration to OSA-Express. The SNA traffic continues to be mission-critical with response time a key consideration.
- Web traffic is growing rapidly and will overtake SNA traffic in volume within the planning horizon.

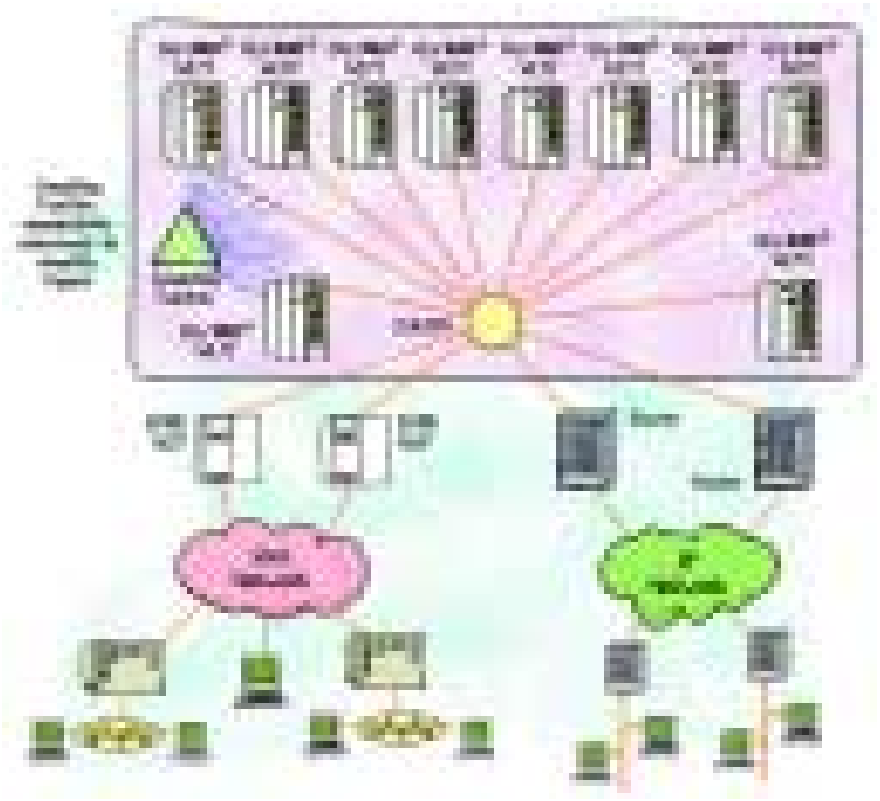

Considerations

- This scenario is similar to Scenario 8 but with one key difference, which is the introduction of the Parallel Sysplex into the data center. This has introduced the requirement for APPN and HPR on S/390 servers, and also, the new processors have OSA installed.
- As in Scenario 8, there are a number of choices for transporting the SNA traffic over the IP backbone network, including:
- Enterprise Extender (EE) to carry APPN/HPR traffic over the IP backbone using UDP
- TN3270E Server, where appropriate
- Data link switching (DLSw) to encapsulate the SNA in IP packets
- Frame-relay BAN shared data transport using frame-relay, X.25, or ATM networks
- In this scenario, to optimize SNA response time and availability, EE is the preferred technique to transport SNA over the IP backbone. S/390 G5 or G6 are already installed, with OSA-Express in plan. The OS/390 software will already be at a level that supports Enterprise Extender. The result is an IP transport network from the branch all the way to the S/390 with nondisruptive rerouting for network outages. In addition, the router network is enabled to differentiate and prioritize the SNA traffic (interactive SNA verses batch) from the IP traffic. See Scenario 8 and the discussion of EE in the "Technology Overview" section for additional information.
- The TN3270E discussion in Scenario 8 is also appropriate in this scenario. TN3270E on the S/390 server is a strategic option. However, the deployment of TN3270E clients is a key consideration.

- DLSw and BAN are recommended for consideration only in cases where branch routers can not be upgraded to support EE. In this case, Communications Sever for Windows NT should be considered for providing EE function at remote sites.
- Because the Parallel Sysplex has OSA installed, the rapid TCP/IP traffic growth can be met with OSA-Express connected to the switched campus infrastructure.
- The migration to OSA-Express will mean that Network Dispatcher will not be effective at workload balancing when using multiple S/390 logial partitions as discussed earlier. In the OSA-Express environment, the optimum workload-balancing solution is provided by using the Sysplex Distributor function (to be available in 2H 2000). Because workload balancing is a fundamental requirement, it is recommended that the OSA-Express migration not take place until the Sysplex Distributor is installed.

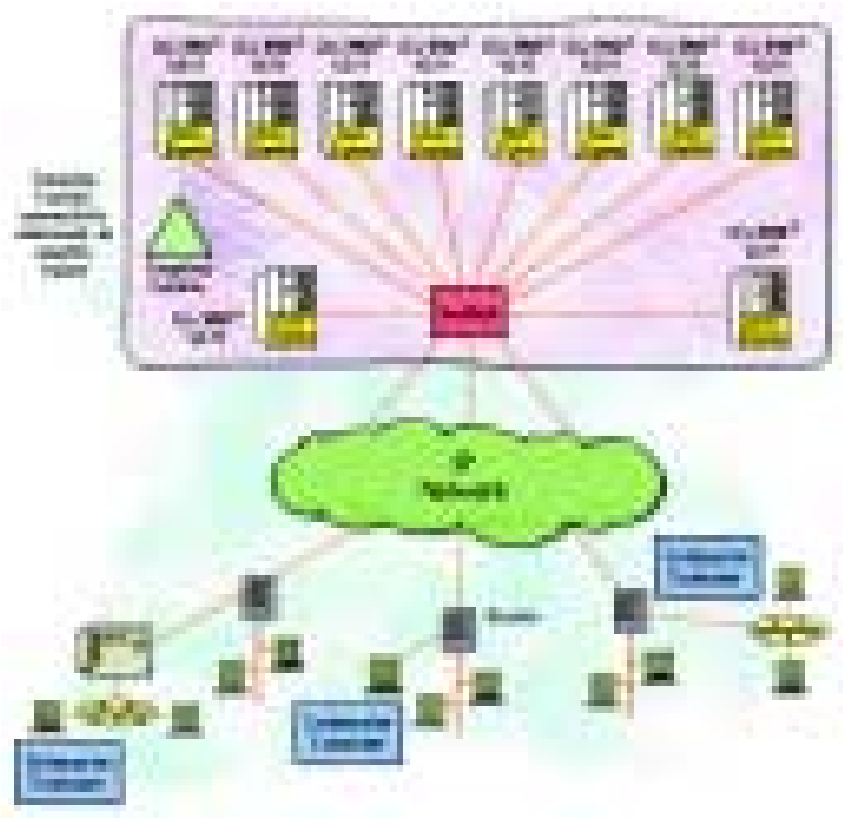

Recommendations

SNA Traffic:

- If TN3270E becomes an option, migrate towards a total IP backbone with TN3270E support provided by the OS/390 software.
- Establish plans to move subarea SNA traffic to an IP transport network. The cornerstone will be EE in the branch and on the S/390 provided by VTAM. This enables an IP infrastructure from the branch all the way to the S/390. Traffic that previously accessed the S/390 as SNA now utilizes IP access to S/390. EE in the branch is available on Communications Server (OS/2® or Windows NT®) platforms from IBM or with SNASw from Cisco.

Note: For SNA traffic that is not terminated on the S/390 server as IP, SNA access to the S/390 remains a reality and a fundamental requirement for access to the mission-critical applications. For subarea traffic, maintain NCP footprints for boundary function and routing. Consolidate NCPs as workload demands permit. Note that NCP may still be required to support SNI and NCP Extended Options (NEO) discussed in Scenarios 1 and 5.

TCP Traffic:

Migrate the TCP/IP traffic from the existing routers to connect directly to OSA-Express using Layer 3 switches.

Scenario 10:
Access to Transaction Processing Facility (TPF)

Current Environment

- The 3746 is used to provide SNA access to the Transaction Processing Facility (TPF).
- The 3172 provides TCP Offload support for TPF today. (In some environments, 3172 is already migrated to a Cisco router with a CIP card.)
- Currently, both the SNA and TCP networks are connected via token ring.

Current Plans

- The 3172 must be replaced because hardware can no longer be ordered and TCP network access is growing.
- TCP Offload performance can no longer handle the throughput requirements of the growing TCP traffic.

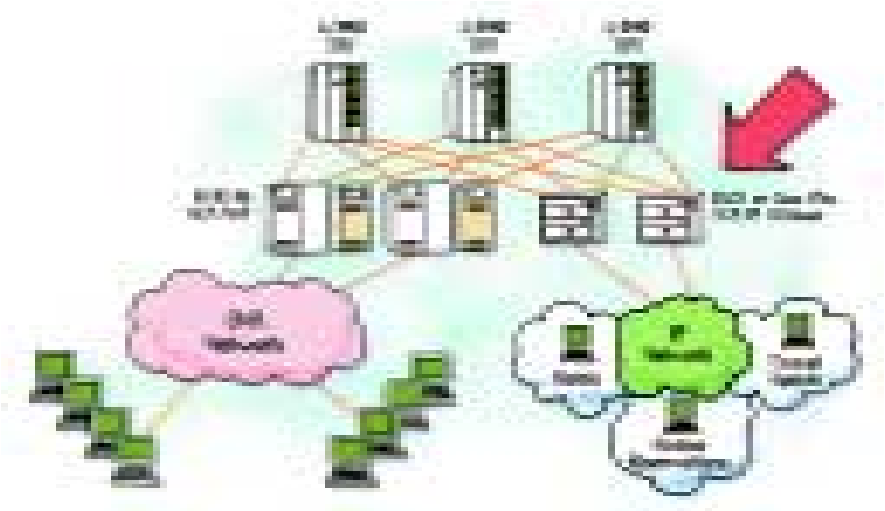

Considerations

- With the latest TPF software release (Program Update Tape 11), a full TCP/IP stack is now provided for the first time. This new TCP/IP stack means that TCP Offload is no longer required for TCP access to IP. Instead, native IP support is provided using the Channel Data Link Control (CDLC) protocol. CDLC support on the 3746 enables the 3746's IP routing to provide a higher-capacity IP connection to TPF natively across the channel. Performance measurements indicate that the new IP native stack on TPF will achieve:
  - TPF throughput (messages per second) for a TCP socket that will at least double, and can be up to 16 times the performance compared with TPF connected to a Cisco 7500 using TCP Offload. The performance improvements when migrating from a 3172 will be even more significant.
  - In addition, that the TPF native stack uses 9 to 17 times fewer CPU cycles in the TPF server, compared to the number of CPU cycles in the TPF server host running TCP Offload.

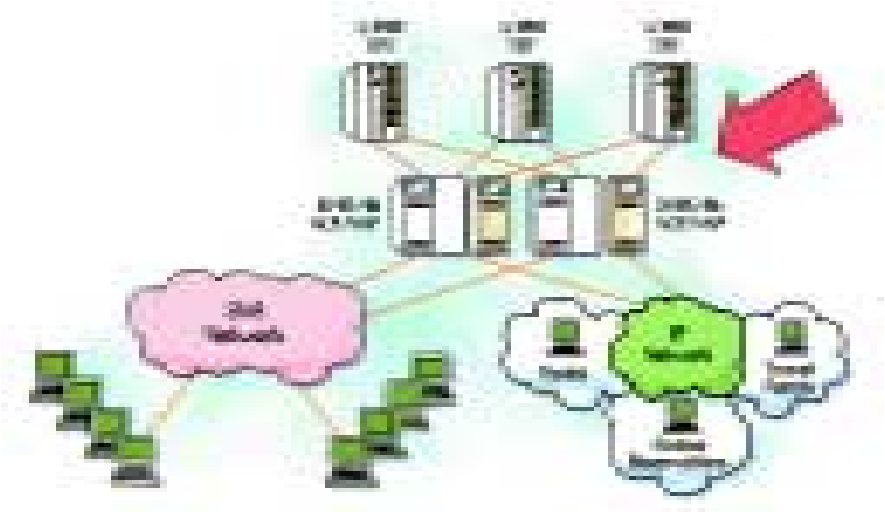

Recommendations

SNA Traffic:

- Continue to use the 3746 for SNA access.

TCP Traffic:

- Upgrade TPF with latest software (PUT11).
- Move the 3172 token-ring connections (and, if installed, the Cisco ones) onto 3746. Upgrade the 3746 with the IP routing feature, and additional processors as required, to handle network growth.

Sources of additional information

Product Information

For general IBM networking information go to:
ibm.com/networking/

For the expanded version of this article "S/390 Server Access Solutions—3745/3746 Evolution Strategy", go to:
ibm.com/networking/nethard.html

For information on the 3745 and NCP, go to:
ibm.com/networking/3745

For information on the 3746, go to:
ibm.com/networking/3746

For information on the OS/390 operating system, go to:
ibm.com/s390/os390/

White Papers

"Optimizing Your 1999 Investment in IBM S/390 Server Gateways: Year 2000 Considerations" at:
ibm.com/networking/nethard.html

"Networking Implications of S/390 Parallel Sysplex" at:
ibm.com/networking/nethard.html

"Host Utilization in Front End Processors: Estimates for the Difference in Host Utilization and Storage between 3745 and 3746 controllers and a Cisco CIP Router" at:
ibm.com/software/network/commserver/library/whitepapers/white_vtamhost.html

John Foley, Jr. (foleyjoh@us.ibm.com) has been a member of the NCP organization for 10 years. His responsibilities include worldwide marketing and sales of the 3745, 3746, NCP and its associated family of products.

IBM Enterprise Storage Server

by Kim Frank

ESS provides:

- ✓ *Ultra-high availability*
- ✓ *Massive scalability 420 GB to 11.2 TB*
- ✓ *High performance with I/O rates exceeding 41000 per second*
- ✓ *SAN-ready*
- ✓ *Advanced copy features*
- ✓ *Log structured array storage options*
- ✓ *Heterogeneous support for Windows NT®, UNIX®, AS/400® and S/390® data storage*
- ✓ *Business-continuance options*
- ✓ *Investment protection*
- ✓ *Flexible IBM lease terms*
- ✓ *Price-competitive0*
- ✓ *Non-disruptive upgrade options*
- ✓ *RAID-5 or non-RAID*
- ✓ *Automatic I/O load balancing*
- ✓ *Synergy with IBM tape devices*
- ✓ *SCSI, Fibre Channel, ESCON®, FICON™*
- ✓ *Full hardware redundancy*
- ✓ *Significant new S/390 ESS EX performance package*

The Storage Server Standard for the New Millennium

e-business is driving a data explosion, generating exponential growth in the need for immediately accessible, always-available, and highly functional storage capacity. In the increasingly e-centric world, information demand is like electricity: you plug in, you get it. Appliance-like intelligent storage served by its own network is required to provide "information as a utility."

The IBM Enterprise Storage Server™ (ESS) is the ultimate storage area network (SAN) utility, providing the information "fuel" that runs the e-business "engine." Extensive heterogeneous server connectivity makes the ESS a natural fit for server-consolidation requirements. The ESS supports rapid universal access to vast quantities of data through many advanced functions and features, making it a workhorse for support of business intelligence and other business-critical applications. Because of its enterprise-wide support and management scope, the ESS is tailor-made to help provide consistent, efficient and effective enterprise resource planning. Using the ESS to address any or all of your strategic and tactical business initiatives will give your organization the business advantage needed to survive and thrive in the e-world.

The ESS is the third generation of the IBM Seascape™ architecture for disk systems. It is a solution that provides the outboard intelligence required by SAN solutions, offloading key functions from host servers and freeing up valuable processing power for applications. As a comprehensive SAN-based storage solution, the ESS provides considerable management flexibility to meet the fast-paced requirements of today's businesses.

- ESS allows you to consolidate your d*ata from different platforms on a single high-performance, high-availability storage server. For example, 50 Windows NT® and 5 UNIX® servers could use Fibre Channel connected to the ESS.*

- *The integrated Web-based StorWatch ESS Specialist is used to configure the ESS subsystem.*

- *The optional StorWatch ESS Expert product provides asset, capacity and performance management that is especially useful for open systems. The Expert provides data such as number of I/O requests, bytes transferred, and read and write response times.*

- *With a capacity of up to 11 TB and up to 32 direct host connections (more with daisy-chaining and switches), an ESS can meet both your high-capacity requirements and your performance expectations.*

- *Most upgrades such as additional capacity are non-disruptive.*

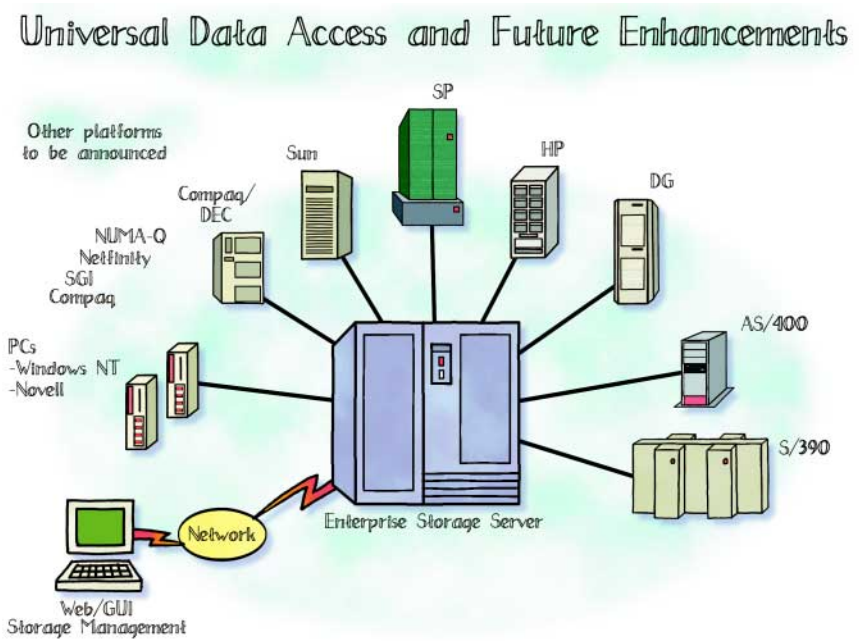

Universal Data Access and Future Enhancements

The ESS supports many diverse platforms including the IBM AIX® operating system, many leading UNIX® variants, Microsoft® Windows NT® and Novell NetWare running on IBM Netfinity® servers and other Intel® processor-based PC servers, and OS/400® software running on AS/400® servers. In addition, the ESS supports System/390® servers running OS/390®, VM/ESA®, VSE/ESA® and TPF software. This rich support of server platforms is not limited in any way; any combination of these heterogeneous platforms can be used with the ESS. Storage capacity is partitioned among the attached servers using the flexible, Web-based StorWatch™ Enterprise Storage Server Specialist™ management tool. The ESSNet feature even provides a browser for the StorWatch interface.

Different platform types use different connection protocols, and the ESS is initially equipped to handle either SCSI (20 MB or 40 MB) connectivity for UNIX, Intel-based PC and AS/400 servers, or ESCON® software for S/390 servers. Fibre Channel is initially supported by the IBM SAN Data Gateway, which provides support for Fibre Channel attachment to the ESS SCSI ports. IBM is previewing plans for the ESS that provide a basis for future development of full SAN exploitation in areas such as disk pooling, file pooling, and copy services. Up to 16 Fibre Channel ports are planned to be available on an ESS. Fibre Channel ports are planned to support FICON™, the Fibre Channel interface for S/390 servers. Direct Fibre Channel attachment is available for short-wave connections to Windows NT, Novell NetWare, and UNIX servers from IBM, Sun Microsystems and Hewlett-Packard.

IBM plans continuing growth and enhancements to the ESS and to make these enhancements available as upgrades to installed ESSs. Planned enhancements such as the implementation of a virtual architecture similar to the RAMAC® Virtual Array (RVA) will triple the current maximum storage capacity, and provide continued performance enhancements by utilizing faster RISC microprocessor engines.

SCSI, FICON, Fibre Channel, and ESCON provide robust host-server connectivity.

PPRC, XRC, Concurrent copy and FlashCopy provide data-copy functions.

Maintenance and upgrades can normally be performed concurrently with full operation, including the concurrent activation of microcode.

ESS "future-proofs" users by providing a new "snap-in" design. With ESS, you can easily add performance, capacity and connectivity upgrades, providing investment protection and 24x7 productivity, saving money and time.

ESS has one of the smallest footprints in the industry for 11-TB configurations.

Business Value of the ESS

In today's information-dependent world of e-business, the ability to gather and store information from every imaginable source and to deliver it to users across the extended enterprise is the key to business success. The ESS provides:

- *Extensive connectivity*: The ESS provides simultaneous attachment to a wide variety of host servers with storage capacity partitioned among the attached servers using flexible Web-based StorWatch management tools.

- *Copy service functions*: The ESS supports several hardware copy functions for two purposes: mirroring functions for disaster-recovery solutions, and copy functions that provide a near-instant copy of data.

- *Comprehensive availability for applications and e-business 24x7x365*: Everything is designed for redundancy: power, cooling, adapters, buses and microprocessor complexes. Redundancy of components means that ESS has been designed with virtually no single points of failure when it comes to hardware components.

- *Breakthrough performance*: The ESS hardware can be configured to provide the capacity and performance required for applications supporting over 30000 I/Os per second. New OS/390 options allow multiple simultaneous I/O operations to the same logical volume, which reduces I/O system queue (IOSQ) time.

- *Low total cost of ownership*: Use of snap-in building blocks keeps IBM development costs low—with the savings passed on to you. IBM Global Financing can develop a lease to fit every requirement. A three-year warranty is included. A one-year warranty is optional.

- *Investment protection*: Existing IBM 7133 Serial Disk System Models 020 or D40 drawers and Versatile Storage Server capacity can be reused in the ESS, resulting in significant cost savings.

- *Extreme scalability*: The ESS adapts to growing storage requirements created by e-business applications scaling from 420 GB to over 11 TB with non-disruptive upgrades. Also, storage consolidation from multiple platforms reduces costs.

- *e-business*: When you're an e-business, the store is always open, so you need data storage that's available 24x7x365. Your customers drive your workload, so you need high-performance data storage that can meet unexpected demands. Customers can flock to your door without warning, so you need storage that is scalable. Suddenly, everyone in the world is your customer, so you need data storage on which you can bet your business!

- *Storage area networks (SANs)*: The ESS offers you the leading storage subsystem, which is supporting the key servers in the industry and enabling customers to realize the promise of nonproprietary SANs.

- Remote hardware support - The IBM hardware support center "call home" connection results in virtually no user intervention required on hardware failures.

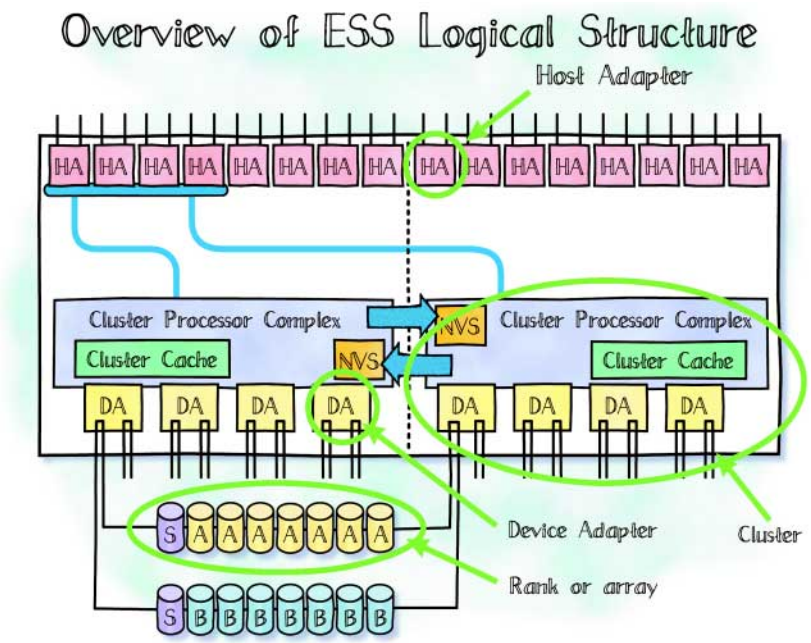

- *ESS has full redundancy of hardware components.*
- *Designed to have the highest quality ever for an IBM disk-storage subsystem.*
- *ESS requires two power cords for redundancy*
- *Some subsystems have only one copy of data in cache—if the cache fails, data is lost.*
- *If a single hardware cluster fails, there should be no loss of data access, because the other cluster has a disk adapter that can access all the data. Write data is retrieved from NVS.*

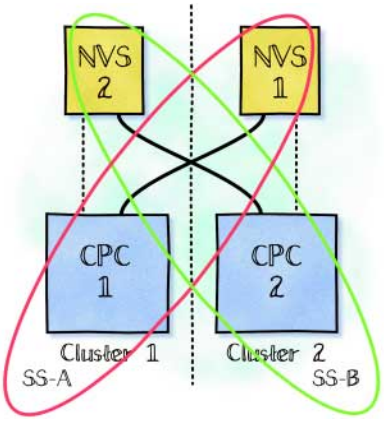

Architecture Overview

The following illustration shows the major components of the ESS. Starting at the top, there are 16 host adapters (HAs), which are the external interfaces. Each adapter supports two ports, for either a SCSI or ESCON connection. Each Fibre Channel HA supports one channel. Each HA connects to *both* Cluster Processor Complexes so that either cluster can handle I/Os from any host adapter.

There are two Cluster Processor Complexes in the ESS, each working independently. Each contains four 64-bit RISC processors with up to 8 GB of cache, nonvolatile storage (NVS) cache, and four device adapters. All host data is written to and from the cache using adaptive-record, partial-track, sequential or full-track staging algorithms. In addition, all write data is automatically placed into two separate caches for integrity. The NVS cache has its own 7-day battery and is managed using a least recently used (LRU) algorithm.

Overview of ESS Logical Structure

On data writes, one copy of data is placed in one cluster and the second copy of the write data is placed in the NVS of the other cluster. In the event of either a planned or unplanned cluster outage, write data for the failed cluster is available on the other cluster and the remaining cluster takes over the functions in progress from the failed cluster. The design philosophy is that once data is in the ESS, it is intended to be fully protected.

A RAID-5 collection of disks is called a rank. The figure shows six RAID ranks labeled A to F. A and B both have spare disks that are used across the loop in case of a disk failure. In the event of a disk failure, the failed disk is replaced and becomes the new spare. The rank is the basic unit of capacity in an ESS and is owned by one logical subsystem, in either fixed-block (open systems) or count key data (S/390) format. A rank is formatted as a set of logical volumes (LVs). The number of LVs in a rank depends on the capacity of the disks in the rank (for example 9 GB, 18 GB or 36 GB) and the capacity of the logical disks being emulated (for example, 3390-3 for a S/390 server).

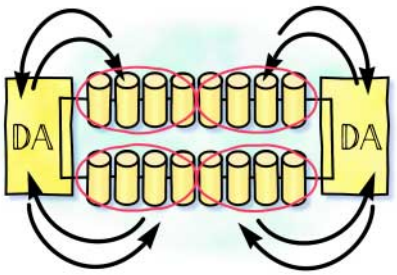

Up to eight simultaneous read/write I/O operations can be performed at once per loop. Up to eight loops and 64 concurrent I/O operations can be performed.

- *Use 18-GB drives for most applications.*
- *Use 9-GB drives for the most demanding performance requirements.*
- *Use 36-GB drives for large capacity and less demanding performance requirements.*
- *Every disk can transfer data at the same time*

Device Adapters

Within each cluster, the device adapters (DAs) are used to connect disks to the Cluster Processor Complexes. DAs are always installed in pairs, one in each cluster. Disk arrays, or ranks (IBM 9-GB, 18-GB or 36-GB disks) are attached to the two DAs via 160-MB Serial Storage Architecture (SSA) loops. The ranks can be configured as RAID-5 or non-RAID arrays ("Just a Bunch of Disks" [JBOD]).

The ESS uses the latest IBM 160-MB serial storage technology in its device adapters, which connect disks to the cluster processors. Disks are connected through a serial-loop architecture with two simultaneous read and write data links. Each link operates at 40 MBps for a total bandwidth of 160 MBps around the loop. Each device adapter supports two independent SSA loops, giving a total bandwidth of 320 MBps. There are four pairs of device adapters in an ESS, providing a total disk-bandwidth capability of 1.280 GBps.

A RAID-5 collection of disks is called a *rank*. The figure shows six RAID ranks labeled A to F. A and B both have spare disks that are used across the loop in case of a disk failure. In the event of a disk failure, the failed disk is replaced and becomes the new spare. The rank is the basic unit of capacity in an ESS and is owned by one logical subsystem, in either fixed-block (open systems) or count key data (S/390) format. A rank is formatted as a set of logical volumes (LVs). The number of LVs in a rank depends on the capacity of the disks in the rank (for example 9 GB, 18 GB or 36 GB) and the capacity of the logical disks being emulated (for example, 3390-3 for a S/390 server).

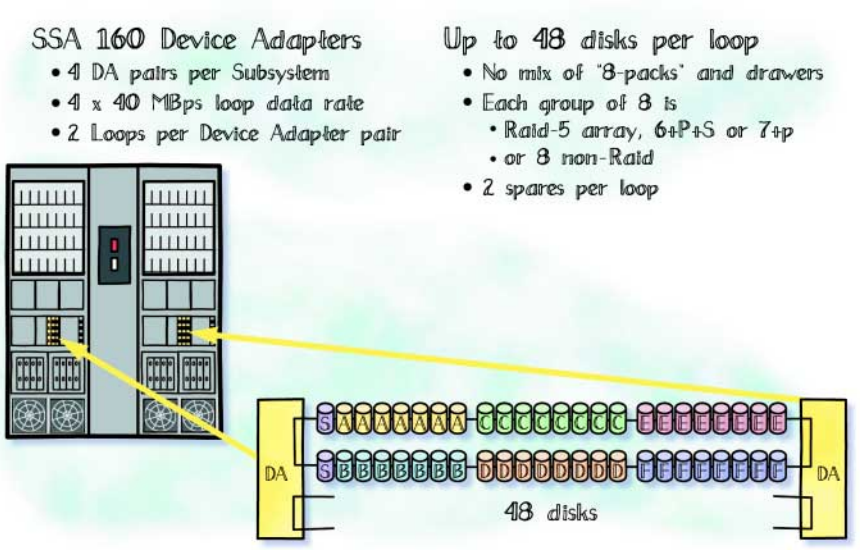

- *Up to 128 disks can be installed in the ESS base frame. For larger-capacity requirements, the ESS Expansion enclosure is available, which holds an additional 256 disks.*
- *Due to advanced designs, there is negligible RAID-5 write penalty.*
- *Data is striped across RAID ranks, which achieves the highest-possible device performance.*

As an example, a group of eight disks (6+P+S) has the capacity of about 53 GB (assuming 9-GB disks). This can be formatted into eighteen 3390-3 disks or one 53-GB logical unit number (LUN) for use by a UNIX system. RAID operations (such as parity generation, and so on) are managed by the DA. A non-RAID option is also available. However, in the event of a disk failure, all data on the failed disk would be lost. This option might be used if mirroring is done by the server's operating system or application software.

You have a choice of 32 standard configurations to meet your business needs. Capacity on the standard configuration scales from 420 GB to over 11 TB. Portions of the storage can be configured as RAID-5 or non-RAID. Storage capacity is customized as logical volumes of variable sizes. LUNs are used for UNIX, Windows NT and AS/400 storage. S/390 storage can be identified as 3880 or 3390 track format volumes.

- *ESS has internal batteries that supply power for a few minutes when power is lost, or to bridge during fluctuations.*
- *For extended power outages the ESS will shutdown gracefully after destaging of all write data.*
- *Many configurations are possible. The choice depends on your performance and capacity requirements.*
- *Many new hardware functions that were not possible before are now available for open systems hosts.*
- *For the S/390 server, there are new performance features that can significantly reduce execution run times.*

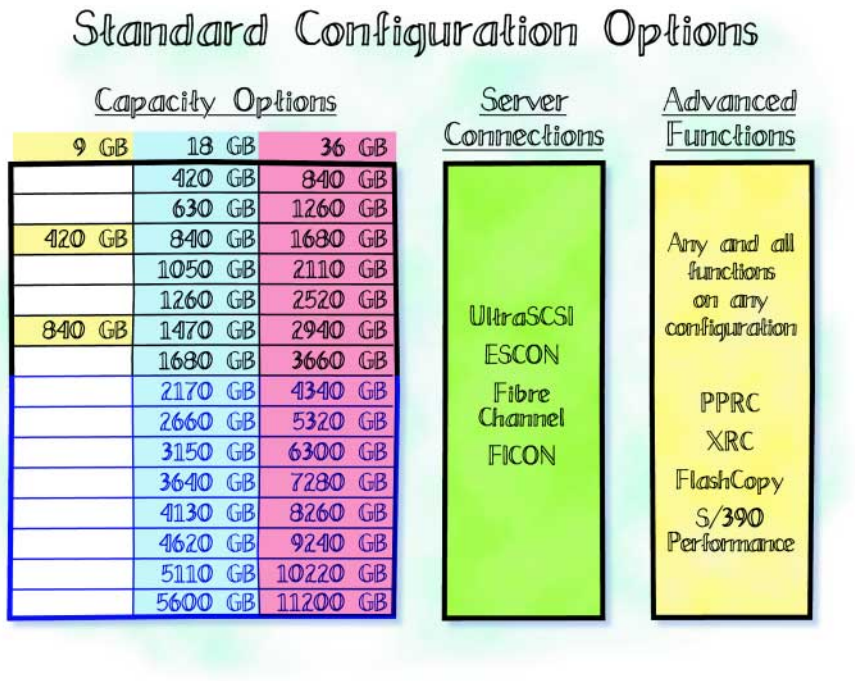

There are various configuration options depending on the operating systems you plan to use on the ESS. Here are some of the choices:

Open System Considerations

- Open system support
  - RS/6000®, AIX
  - PC Server, Windows NT 4.0
  - AS/400, OS/400
  - Sun Solaris
  - HP 9000, HP-UX
  - Data General, DG-UX
- UNIX Systems SCSI ID, LUN
  - 15 targets, 64 LUNs
  - 0.5 GB, 226-GB logical volume
- AS/400 with OS/400
  - 9337-580/590/5AC/5BC
  - 4.19 GB, 8.59 GB, 16 GB, 32 GB
- Data Path Optimizer for AIX, Windows NT
- Configure using ESS Specialist

S/390 Considerations

- 16 x 3990 CU images
  - 3990-6 emulation
- Up to 256 devices/CU image
  - 3390-2/3/9 with 3390 Track
  - 3390-2/3 with 3380 Track
- Maximum of 32 ESCON ports
  - Each with 64 logical paths
  - Maximum of 2048 logical paths
  - 128 logical paths per CU image
- Configure using ESS Specialist
- OS/390, VM, VSE, TPF

- *ESS supports AS/400®, RS/6000®, Netfinity®, NUMA-Q, Intel®, HP, Sun and other host servers. For a complete list, see:* ***ibm.com****/storage/hardsoft/products/ess/supserver.htm*
- *ESS offers open systems users 24x7 availability with redundant design and non-disruptive maintenance.*
- *The StorWatch ESS Specialist enables centralized, physical storage management, which reduces cost and management complexity.*
- *IBM 7133 and VSS storage systems capacity can be used with the ESS, providing investment protection for assets already paid for.*
- *TSM (formerly ADSM) can be used to back up ESS storage.*

ESS Open System Overview

The ESS provides universal data access including: RS/6000 and SP2® running AIX, many leading UNIX variants, IBM Netfinity and other Intel processor-based PC servers running Windows NT and Novell NetWare, and AS/400 running OS/400. This rich support of server platforms is not limited in any way; any combination of these heterogeneous platforms can be used with the ESS (including S/390). The storage capacity is partitioned among the attached servers using flexible Web-based StorWatch Enterprise Storage Server Specialist management tools. ESSs in one location are attached to a LAN and to an external processor (supplied with the ESS) that provides a means of accessing the StorWatch tools and can provide a gateway to your WAN for remote access and management from anywhere in the world.

IBM StorWatch Enterprise Storage Specialist

StorWatch ESS Specialist is used to configure, install and allocate storage to the various host servers that will use the ESS.

For SCSI-attached hosts, a storage facility consists of one or more SCSI I/O interfaces (buses), each with one or more SCSI targets. Each target has an associated set of logical units. Each logical unit has a unique LUN on the target. A given target is accessible through a single SCSI bus. A given SCSI bus can access one or more targets. A set of logical devices is associated with a logical subsystem. A logical subsystem is managed by a single device adapter and a single Cluster Processor Complex at any given time. The ESS has two Cluster Processor Complexes, so if one fails or is taken offline, the other cluster takes over all the work on the cluster that is no longer available. When one cluster is taken offline, write data that has not been destaged is retrieved from the other cluster's NVS. Normal operations continue.

The figure shows a sample panel.

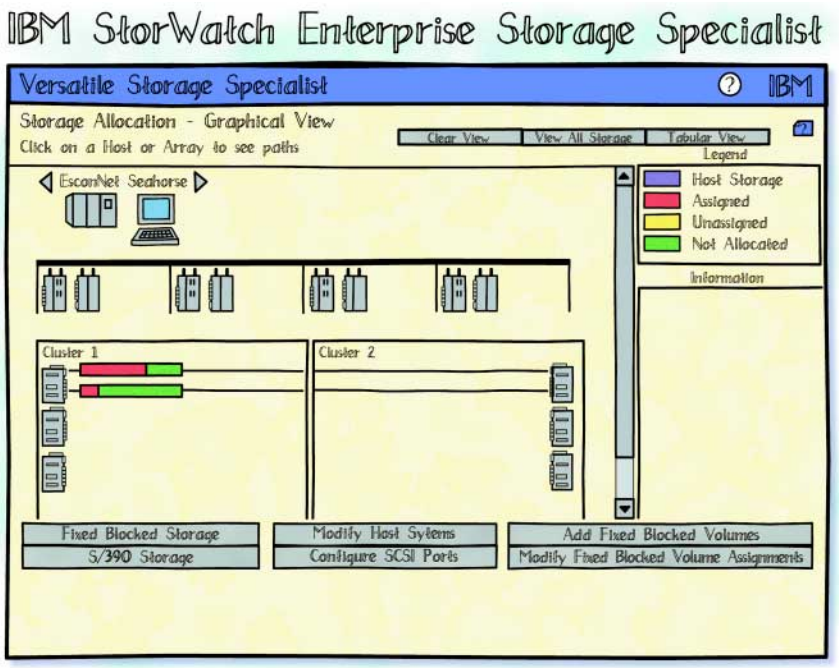

If you imagine a SCSI host's view of an ESS, it looks like a group of SCSI disks attached to a SCSI bus. The actual number that any UNIX or Windows NT system can support is considerably less than the maximum shown in the following figure.

- *Each fixed-block (FB) logical subsystem supports up to 256 logical volumes.*
- *Logical volume size can vary from 0.5 GB to 245 GB (the size of a RAID rank with 36-GB disks).*
- *A total of 4096 logical devices can be defined.*

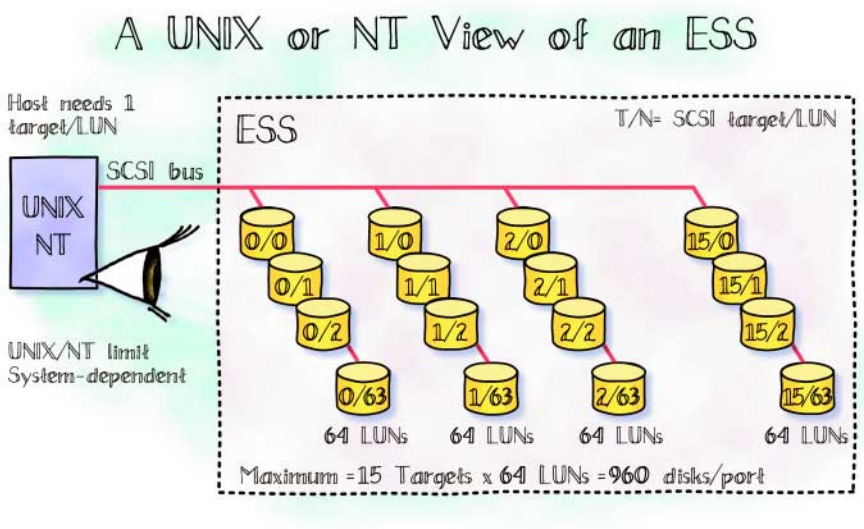

IBM Subsystem Device Driver (SDD) provides:

- *More than one path from the host to the ESS*
- *A single LUN can appear as 2 to 16 LUNs*
- *Host path failover*
- *Load distribution across paths*
- *Support for SCSI AIX® and Windows NT*
- *Support for Windows NT, Windows® 2000 and AIX Fibre Channel*
- *Sun™ and HP native Fibre Channel*
- *HACMP support is planned for a future release*

One target/LUN is used for each host attached to the ESS SCSI port. Typically, you will have only one host per SCSI port, leaving 15 target IDs and a number of LUNs per target. For availability purposes, you can configure a logical device in the ESS as a shared device using two different SCSI ports in the ESS. This will allow you to interconnect your host to two or more separate SCSI ports.

If your host OS is AIX or Windows NT, you can use the IBM Subsystem Device Driver (SDD) to distribute I/O activity across the SCSI adapters in the host. SDD will recover I/Os that fail due to a path problem by using an alternative path. A sample configuration is shown in the following figure.

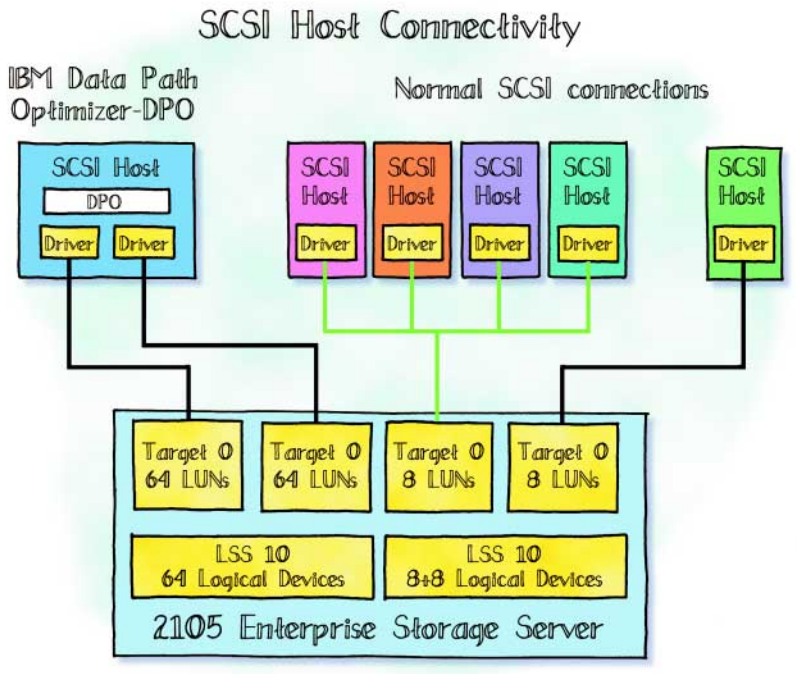

This outstanding performance provides:

- *Better handling of unpredictable e-business workloads*
- *Less manual tuning because the ESS stripes data requests over many disks*
- *Ability to handle peak loads with ease*
- *Ability to run both open systems and S/390 data on the same storage server*
- *Ability to scale much higher than the 4 SCSI interfaces used for this comparison (up to 32 concurrent transfers)*
- *ESS throughput that scales proportionately as you increase the number of host interfaces and capacity*

Open Systems Performance Examples

IBM performance tests have demonstrated that the ESS has significantly exceeded IBM's design goals. The following example used four RS/6000 processors, 32 Ultra-SCSI ports, 128 drives (18.2 GB each) with 4000 records achieving over 40000 I/Os per second. No other competitor came close in those tests.

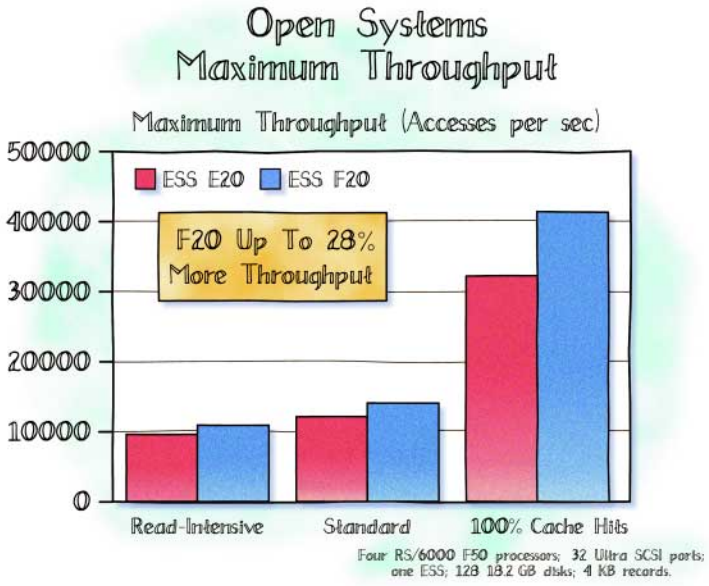

The IBM 7133 family is regarded in the storage industry as the fastest non-cached storage solution (see the following figure). The ESS design enables speeds almost as high as with the 7133, with IBM leaving one storage competitor way behind in read-intensive operations

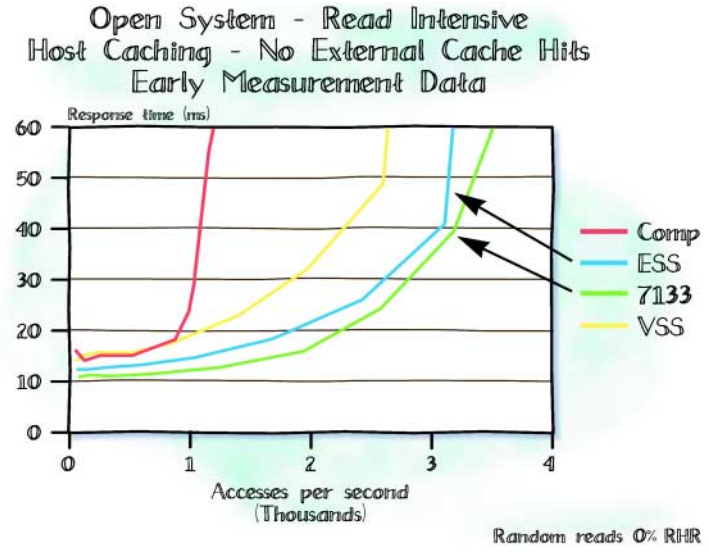

- *ESS hardware provides a number of distinct ways to share and copy open systems data.*
- *PPRC is the ideal way to provide for business continuance in the event of a major outage at one location.*
- *With PPRC, the ESS ensures that write data is in two separate ESSs before letting the application continue.*
- *PPRC is an industry standard.*
- *PPRC is established on a volume level and can be used by midrange and S/390 systems.*
- *PPRC uses direct connections between ESS systems with no host involvement.*
- *FlashCopy is designed to make copies of volumes available in less than 2 seconds.*
- *PPRC and FlashCopy are separate options on ESS systems.*
- *More details about PPRC and FlashCopy are in the "S/390 Overview" section.*

Open Systems Copy Functions

Some of the advantages with an ESS in an open systems environment are the excellent performance, reliability, scalability, and serviceability it offers. In addition to that, the ESS provides new functions for the open systems environment. These include:

- Data Sharing
  - Access to the same data from like systems
  - Host software must manage data integrity
- Peer-to-Peer Remote Copy
  - Disaster-recovery solution
- FlashCopy
  - Immediate copy available for backup, data mining and other use

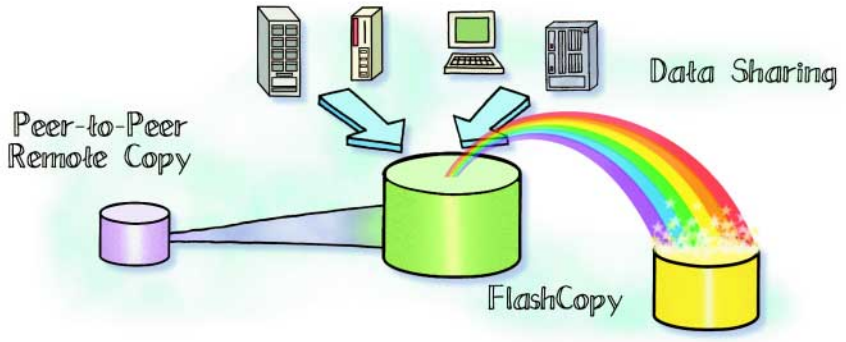

The capability of ESS to share LVs between like systems allows concurrent access to databases by different hosts (both IBM DB2® Universal Database™ and Oracle Parallel Edition support this feature). The database products provide locking mechanisms to guarantee data integrity.

With the Peer-to-Peer Remote Copy (PPRC) function that ESS offers, midrange systems such as Windows NT, NetWare and UNIX (except for the AS/400) can now realize a disaster-recovery solution to protect data should one location be destroyed. PPRC provides a synchronous copy of selected data up to a distance of 103 km (64 miles). PPRC is administered from the Web browser interface of StorWatch ESS Specialist. No special software is required to use PPRC. ESCON cables connect the ESS systems. Write operations are completed on both copies (local and remote) before they are reported as completed to the operating system.

The FlashCopy™ feature enables you to create a copy of your data immediately, whenever you need it. FlashCopy generates an instant copy of a logical disk. For example, FlashCopy could be used to enhance asynchronous backup, creating test data from production data, checkpoints, and for non-disruptive copying for business-intelligence applications. The copy of the data is available immediately after the command is issued.

- ESS supports *OS/390, VM/ESA, TPF and VSE/ESA operating systems.*
- *Up to sixteen 3990 control units can be defined in the ESS.*
- *Custom volumes can also be configured. Custom volumes allow a volume to have any number of cylinders you choose. You would typically define a small number of cylinders to reduce contention for high-activity data sets.*
- *In most cases, however, the PAV function would be a better choice.*
- *The combination of StorWatch Specialist and HCD/IOCP are used to configure the ESS.*

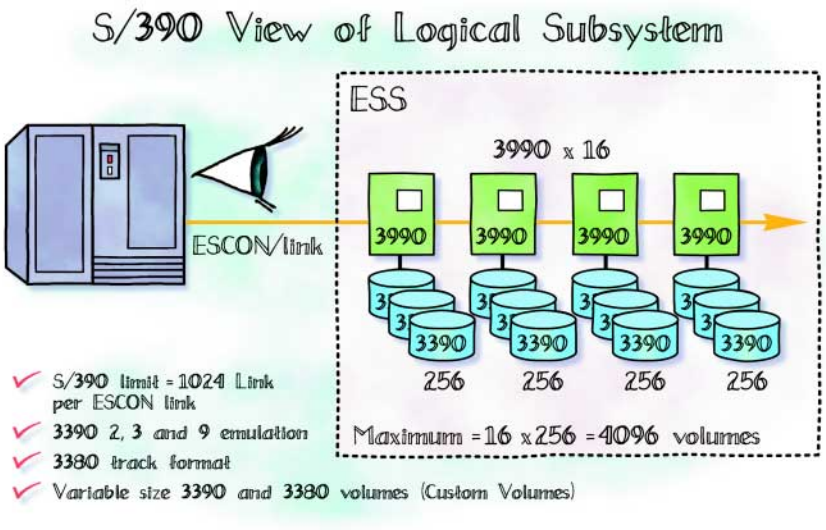

S/390 Overview

The ESS supports S/390 servers and inherits the S/390 functions previously provided by the 3990-6 and RAMAC families. However, the ESS far surpasses these legacy technologies in all areas such as performance, hosts supported, availability, functions, scalability and non-disruptive operation. The ESS is designed specifically to support both

S/390 and open systems workloads at the same time. Typically, S/390 workloads are cache-friendly and take advantage of large caches, whereas the open systems workloads often are very cache-hostile. For the S/390 workloads, there are sophisticated cache-management algorithms and a large cache. For workloads that do not take advantage of large cache sizes, the ESS has high-performance disk arrays with 160-MB disk loops. So whatever the workload—even mixed workloads—the ESS delivers high performance. The S/390 hardware view of the ESS is shown in the following figure.

The number of 3390 or 3380 devices available depends on the type of disk installed (9 GB, 18 GB or 36 GB), the number of disks installed and the type of emulation you select.

IBM has defined standard physical configurations that meet most customers' requirements and are intended to make it easy to install an ESS. The standard configurations include options for ultra-high performance, large capacity, ESCON-only, fixed-block only (open), and mixed environments (SCSI, ESCON). For complex environments, the ESS can be configured to match your exact requirements.

In the following example, two S/390 processors connect via ESCON links to an ESS with eight 3390 control units defined

- *The configuration on the right is fault tolerant in that HAs, clusters, disks, buses, disk adapters, ESCON Directors, or paths can fail and the applications will not be affected.*

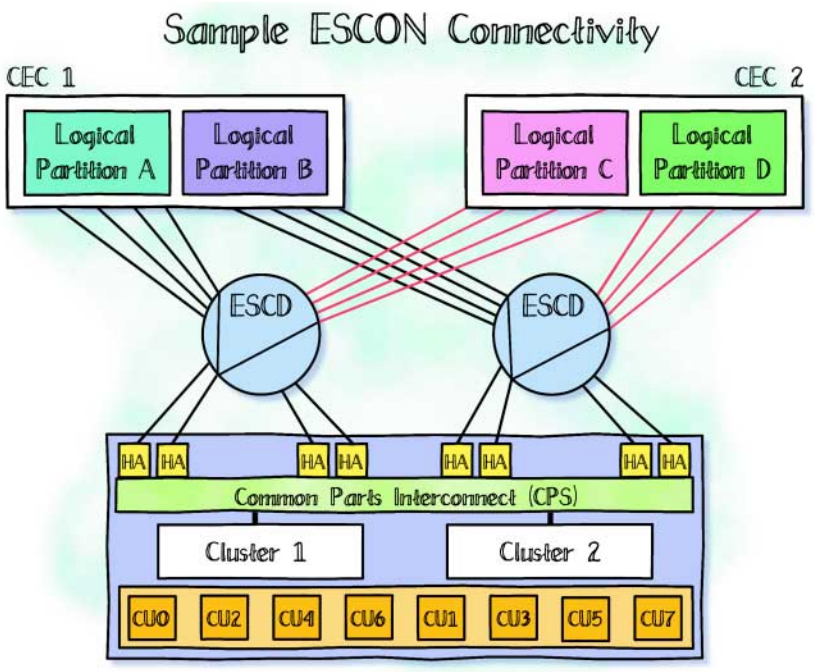

With high-speed microprocessors, 160-MB disk interfaces, large cache and large capacity, the ESS provides extraordinarily high performance. The following illustrations show the performance characteristics of representative workloads that you might have in your environment.

In this example, with a cache-friendly workload, the new ESS Model F20 exceeds the server's capacity to drive it.

- *The chart on the right compares various storage systems from IBM with the new ESS Model F20.*
- *The ESS can handle significantly more work than is shown here as more channels and capacity is added.*

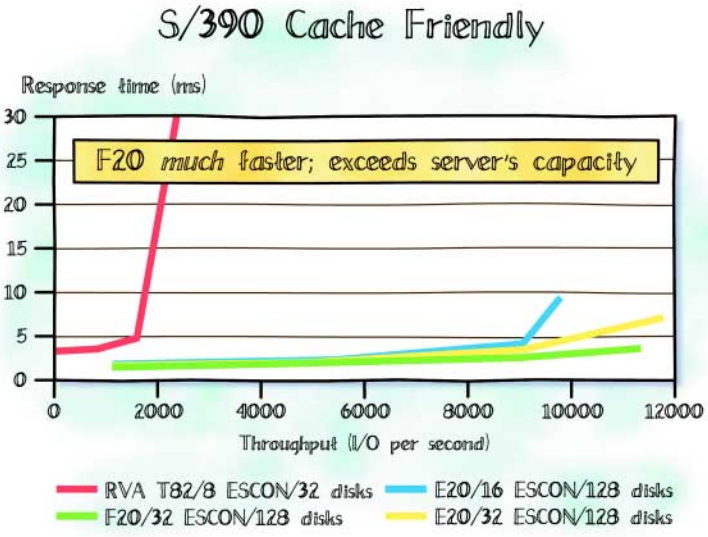

ESS priority I/O queuing:

- *I/Os can be queued in ESS.*
- *Priority can be set by WLM.*
- *I/O priority for systems in a Sysplex is honored.*

Even with cache-hostile workloads, the ESS can drive a very high rate of I/Os with good response times.

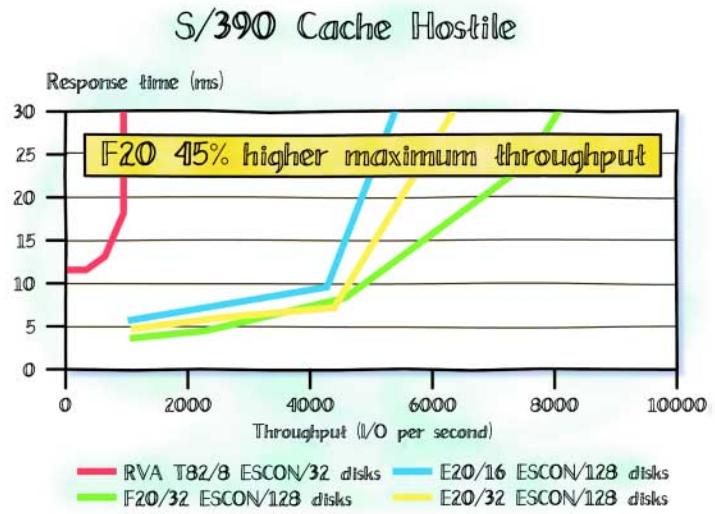

Performance Improvements

The IBM storage server design team knew they would have the fastest storage server in the industry for the S/390 server—but they didn't stop there. They teamed with the IBM OS/390 software development lab and invented three new revolutionary improvements that can increase performance 2 to 3 times over what the raw hardware can deliver. The performance gains you could achieve are entirely dependent on your particular environment. However, as you understand these new innovations, it might become obvious how these innovations could apply to your environment.

These innovations, called the *ESS EX Performance Package*, can dramatically improve performance and significantly reduce the elapsed times of all the applications in your systems:

- *Parallel Access Volumes (PAVs)* allow a single S/390 server to simultaneously process multiple I/O operations to the same logical volume (the SCSI architecture already provides this function for open systems hosts).
- *Multiple Allegiance* is a similar enhancement that expands the simultaneous logical volume access capability across multiple attaching OS/390 servers. These two functions enable the ESS to process more I/Os in parallel, dramatically improving performance and enabling greater use of large volumes, which helps to simplify storage management and to reduce costs.
- *Priority I/O Queuing* allows the ESS to use information provided by the OS/390 Workload Manager (WLM) to manage the processing sequence of I/O operations so that they match application priorities. The ESS has also implemented enhanced commands that significantly reduce overhead, or time spent by the system that is not directly related to the transfer of information, which further improves performance.

S/390 architecture allows only one I/O operation at a time to disk devices. The new ESS changes this so that more than one I/O can be active at a time to one disk device.

OS/390 ESS PAV support includes:

- *Transparency support, which has PTFs to use the ESS as a 3990 only.*
- *Exploitation static support, which has PTFs that recognize the ESS0 with PAV support and can use alias and base addresses for I/O.*
- *Exploitation dynamic support, which uses OS/390 2.7 plus PTFs. The Workload Manager will automatically tune the OS/390 I/O system using the alias address by the use of policies.*

PAVs:

- *Sharply reduce IOSQ time*
- *Provide multiple UCBs per logical volume*
- *Make aliases visible only to I/Os in*

OS/390

- *PAV support is an ESS option.*

Since it was developed, S/390 architecture allowed only one channel program to be active to a disk volume at a time. This ensured that data being accessed by one channel program could not be altered by the activities of another channel program. OS/390 does not try to issue another I/O to a disk volume when a previous I/O is already active for that volume. Unit control blocks (UCBs) and OS/390 queues keep track of I/Os, so they are processed one after another. The ESS together with OS/390 software are changing all this.

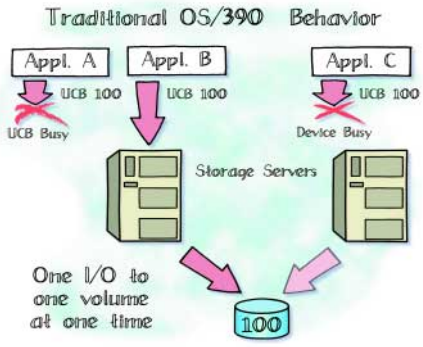

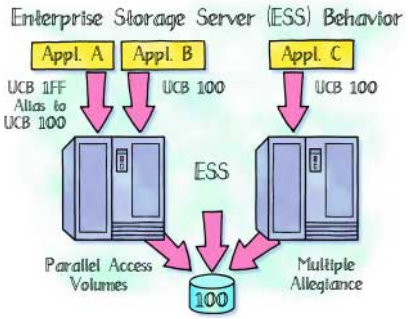

Multiple Allegiance

The ESS introduces the capability to do more than one I/O to a S/390 logical volume. The ESS introduces the concept of *alias addresses*. Instead of one UCB per logical volume, an OS/390 host can now use several UCBs for the same logical volume. The ESS also accepts I/Os to a shared volume coming from different hosts in parallel. This capability is called *Multiple Allegiance*.

In the table on page 58, the first row shows hosts A and B running separately to a DB2® volume. The second row illustrates how host B can significantly slow down host A's I/Os. With the ESS and Multiple Allegiance, there is no slowdown, as shown on the third row.

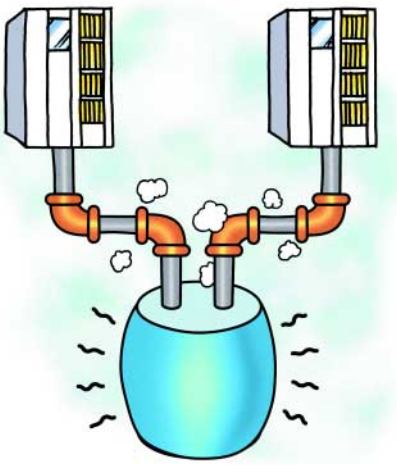

Multiple Allegiance provides:

- Concurrent access from multiple path groups (system images) to a volume
  - Incompatible I/Os are queued in the ESS.
  - Compatible I/O (no extent conflict) can run in parallel.
  - ESS guarantees data integrity.
  - No special host software required. However, host software changes can improve global parallelism (limit extents).
- Improved system throughput
  - Different workloads have less impact on each other.

	Host A (4K read hits)	Host B (32 record 4K read chains)
Max. ops/sec – isolated	767 SIOs/sec	55.1 SIOs/sec
Max. ops/sec – 100% extent conflicts	59.3 SIOs/sec	54.5 SIOs/sec
Max. ops/sec – full Multiple Allegiance	756 SIOs/sec	54.5 SIOs/sec

These state-of-the-art advancements in the ESS result in an extremely fast storage subsystem. Performance is dependent on the workload and particular hardware used. However, the following figure is typical of what you might see when you are using the ESS and PAV.

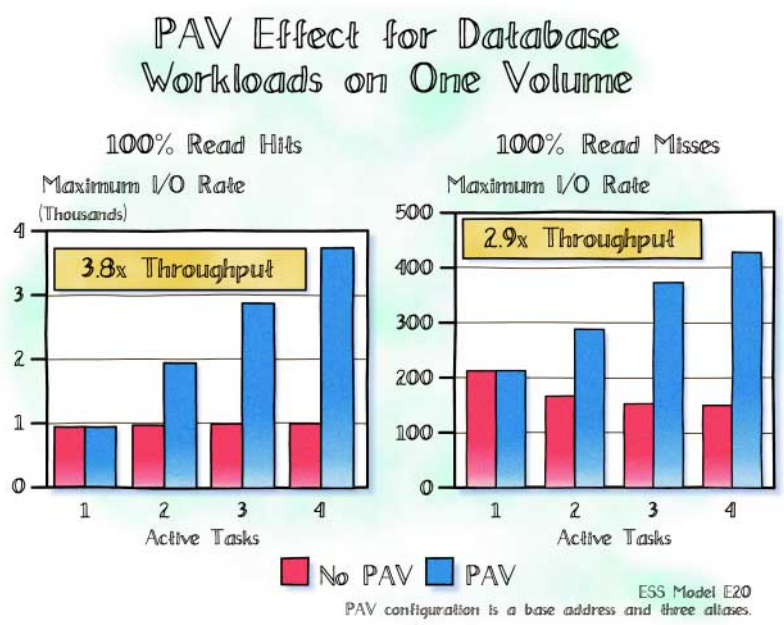

This figure demonstrates the outstanding performance improvements using PAVs. Imagine this effect for more than one volume.

ESS Copy Services

The ESS supports several hardware copy functions for two purposes: mirroring for disaster-recovery solutions, and copy functions that provide an instant copy of data.

- FlashCopy provides an instant copy of a volume for both SCSI and S/390 servers.
- Concurrent Copy for S/390 servers provides a copy of a volume or data set.
- Peer-to-Peer Copy (PPRC) provides synchronous mirrors for SCSI or S/390 servers.
- Extended Remote Copy (XRC) provides asynchronous mirroring for S/390 servers

- *Copy functions are supported and can be initiated by OS/390. UNIX systems use the StorWatch Specialist software to initiate and control copy services.*
- *PPRC and XRC copy architecture is an industry standard.*

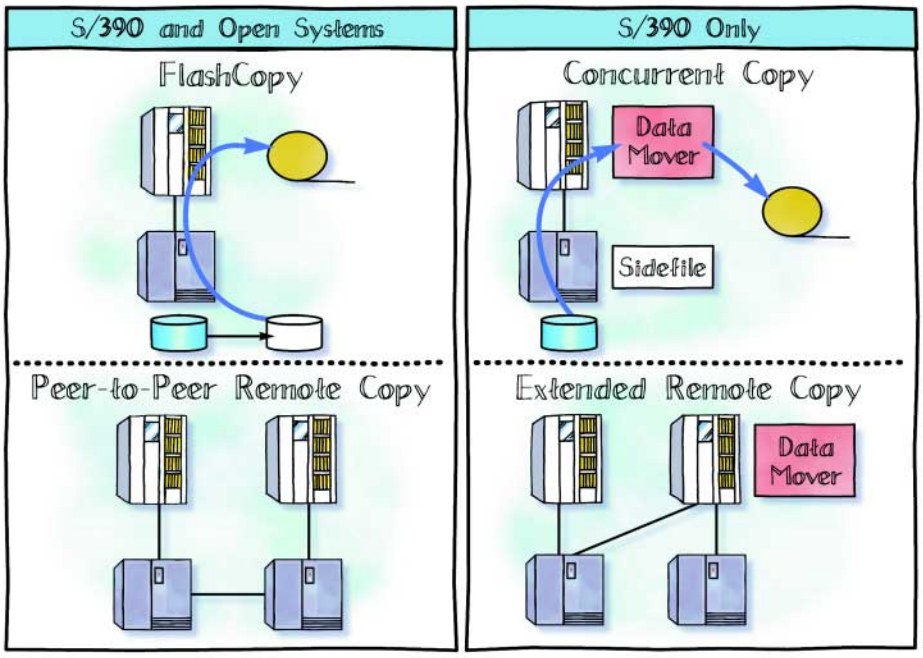

- *FlashCopy is an optional feature of the ESS. Both S/390 and open systems can use it.*
- *FlashCopy can be used to:*
  - *Create test data*
  - *Create temporary checkpoint copies*
  - *Provide asynchronous backup–copy data and make the backup from the copy while your application access the source*
  - *Perform data mining and data warehousing on the copy of the data*
- *FlashCopy can be started using DFSMSdss™ or the StorWatch Specialist.*
- *PPRC is used for business recovery in the event of a data-center disaster.*
- *S/390 and open systems can use the function at the same time.*

FlashCopy and Concurrent Copy

With SNAP/SHOT® on the IBM RAMAC® Virtual Array (RVA) or Concurrent Copy you get a copy of the original volume or data set in a few seconds. The FlashCopy function provides a similar capability on the ESS for volume copies. As with SNAP/SHOT, you get an instant copy when you start the command and it does not require predefined back-end storage.

When the ESS is enhanced with virtual architecture, FlashCopy won't require back-end storage. Once the relationship between the source and target volumes is established, a background task copies the tracks from the source to the target. Because the background task is making the copy, both the source and target volumes can be used for reading or writing independently.

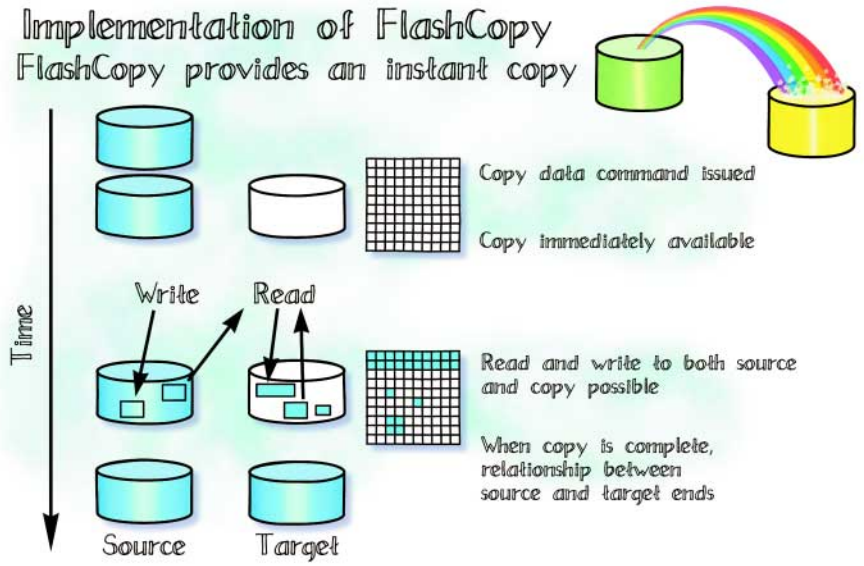

Peer-to-Peer Remote Copy

Concept

- Synchronous copy, mirrored (RAID-1) to another ESS
- Disaster-recovery solution
- Established on a volume level
- Direct connections between ESS systems using ESCON links

Interfaces

- TSO/E for OS/390 systems (and ICKDSF for S/390)
- StorWatch ESS Specialist Web interface for S/390, UNIX and Windows NT

Benefits

- Data-integrity protection
- Remote mirroring up to 100 km (62 miles)
- De facto standard
- Offloads server
- Common facility for all attached hosts

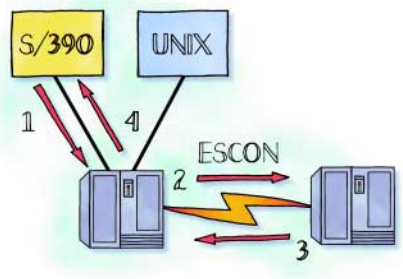

- *PPRC is a defined architecture that has many unique data-integrity and performance features.*
- *PPRC for OS/390 data is controlled by OS/390 ICKDSF or TSO/E commands*
- *Open systems use the StorWatch Specialist Copy Services functions.*
- *Would your critical applications survive a real disaster?*
- *GDPS provides automated services using the ESS to provide near-continuous operation, even in a disaster.*
- *GDPS is only for OS/390 systems.*
- *Offers the combination of CMOS processors, ESS, high performance, and IBM Global Services to provide a comprehensive disaster-recovery solution.*

PPRC is an ESS hardware remote-copy or mirroring solution that can help preserve data integrity, even in a rolling disaster.

PPRC is a synchronous mirroring solution, which means that an I/O does not complete until it is acknowledged from the remote site. PPRC is a *de facto* standard in that other vendors have licensed the architecture from IBM to use in their storage subsystems. PPRC is set up on a volume or LUN basis. Two or more ESSs are connected by ESCON links (fiber-optic links using the ESCON protocol). Updates to a PPRC volume on the primary site go first in the cache and nonvolatile storage in the primary storage control (see (1) in the figure above). The updates are then sent over the ESCON links to the remote ESS (2).

When the data is in cache and NVS on the secondary site, receipt of the data is acknowledged (3) and the primary ESS signals the application that the I/O is complete (4). This process places data in two places before the application considers the operation complete. In the ESS, the efficiency of the ESCON protocol has been significantly improved over prior implementations. It allows for an increase in distances up to 103 km (63 miles).

PPRC is very powerful as a stand-alone solution, but it is not enough for the most-demanding S/390 customers, such as banks, hospitals, e-businesses, and 24x7x365 solutions. If one site goes down, there is no automatic way for the other site to take over or even know a site is no longer available. This is shown in the illustration.

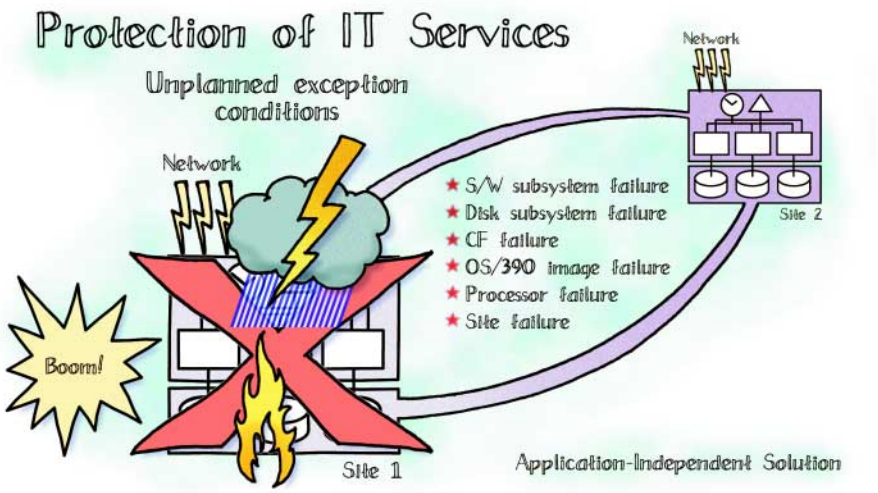

IBM GDPS provides:

- *Near-continuous application and data availability*
- *Planned application availability across both reconfigurations and unplanned reconfigurations*
- *Data consistency and integrity—simplifies operations from a single point of control*
- *Application independence*

A key ingredient of GDPS is the ESS PPRC function, with the high performance and data integrity the ESS provides.

IBM has an answer to this dilemma for S/390. It is called Geographically Dispersed Parallel Sysplex (GDPS). GDPS is a multisite management facility that is a combination of system code and automation. It utilizes the capabilities of Parallel Sysplex technology, storage subsystem mirroring (PPRC) and databases to manage processors, storage and network resources. It is designed to minimize and potentially eliminate the impact of a disaster or planned site outage. It provides the ability to perform a controlled site switch for both planned and unplanned site outages, maintaining full data integrity across multiple volumes and storage subsystems and it enables you to perform a normal DBMS restart (not DBMS recovery) at the opposite site.

IBM has proven the GDPS concept with multiple customers. For example, with GDPS, a recent simulated site disaster at a customer's site caused no data loss and the recovery window was reduced from 12 hours to 22 minutes.

GDPS consists of a base or Parallel Sysplex cluster spread across two sites separated by up to 40 km (25 miles) with one or more OS/390 systems at each site. The multisite Parallel Sysplex cluster must be configured with redundant hardware (for example, a Coupling Facility and a Sysplex Timer® in each site) and the cross-site connections must be redundant.

GDPS consists of production, standby and controlling systems. The production systems execute the mission-critical workload. The standby systems normally run expendable work, which can be displaced to provide processing resources when a production system or a site is unavailable. GDPS uses a combination of storage subsystem, Parallel Sysplex technology, and environmental triggers to capture, at the first indication of a potential disaster, a data-consistent secondary site copy of the data, using the new, recently patented PPRC freeze function.

GDPS consists of on-site services provided by IBM Global Services to plan for the general installation, installation of GDPS automation software, and customization of the environment.

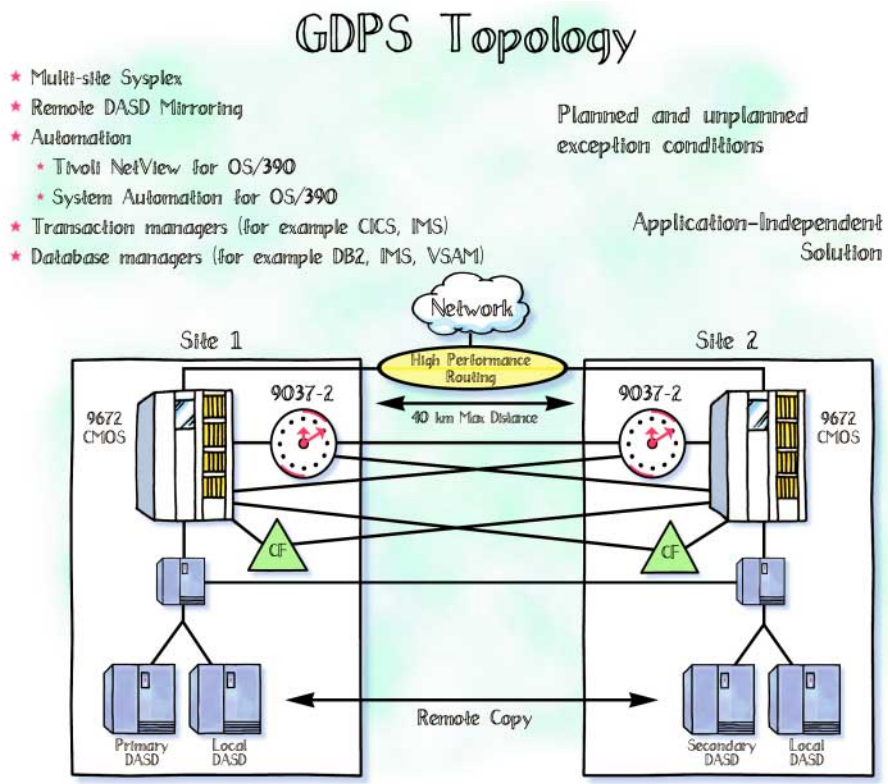

- *Extended Remote Copy(XRC) was previously available on IBM 3990 Control Units and is for OS/390 systems only.*

- *PPRC and XRC can also be used to move data from one ESS to another, for instance if leases are up. P/DAS is an OS/390 function that can be used in conjunction with PPRC and XRC. It is designed to allow primary and secondary volumes to be swapped without an application outage.*

Extended Remote Copy

For those situations that don't require synchronous mirroring or long distances, Extended Remote Copy (XRC) provides the solution. XRC is a de facto standard that is an asynchronous remote-mirroring solution. It requires the OS/390 System Data Mover (SDM) to move the data. Applications doing write I/Os to primary (source) volumes (see 1 in the figure) receive "write complete" as soon as the data has been secured in cache and NVS of the primary ESS (2). The SDM reads the update to XRC source volumes from cache and sends it to the secondary volume on a remote storage subsystem. The ESS implementation of XRC is enhanced with full suspend/resume, which enables the secondary copy to be resynchronized following either a planned or unplanned telecommunications outage without performing a time-consuming full-volume copy operation.

ESS XRC Implementation

- Remote mirorring of volumes
- Asynchronous copy process
- Supports large distances
- The OS/390 System Data Mover (SDM) moves data from primary susystem to a remote storage subsystem
- ESS's XRC implementation is compatible with the previous implementation on IBM 3990-6
- ESS supports all advanced XRC V2 functions:
  - Multiple reader support (max. 64/LSS)
  - Dynamic balancing of application write bandwidth vs. SDM read performance
  - Floating utility device, or use 1-cylinder utility device
- Uses new performance-enhanced CCWs
- XRC for Open Systems volumes possible with XPE
  - Store your Open Systems data on 3390-3 type volumes
  - Copy these volumes with XRC from an OS/390

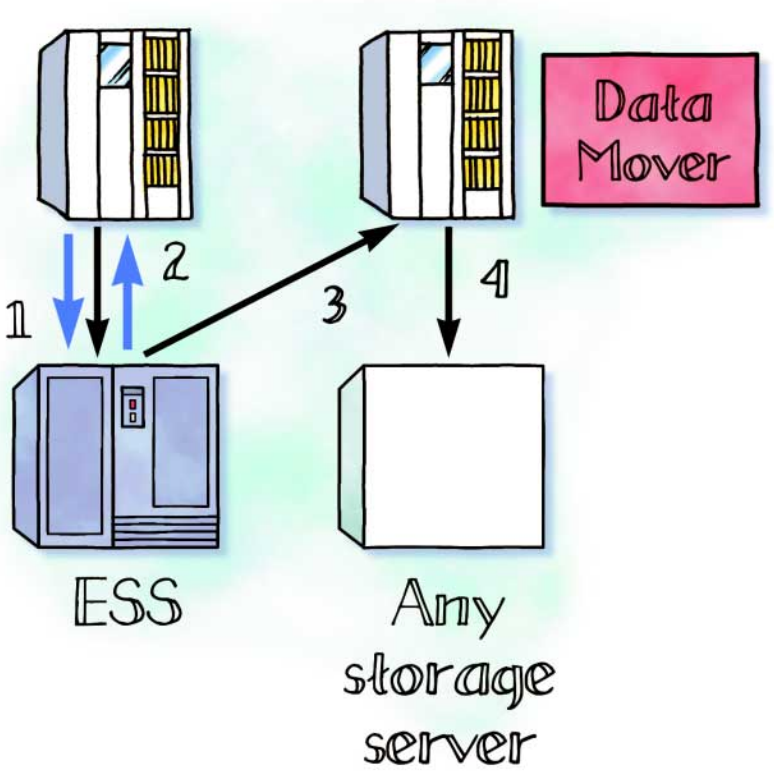

XRC can be used to clone an entire ESS or a subset of volumes. This capability could be used to move a data center from one location to another as shown.

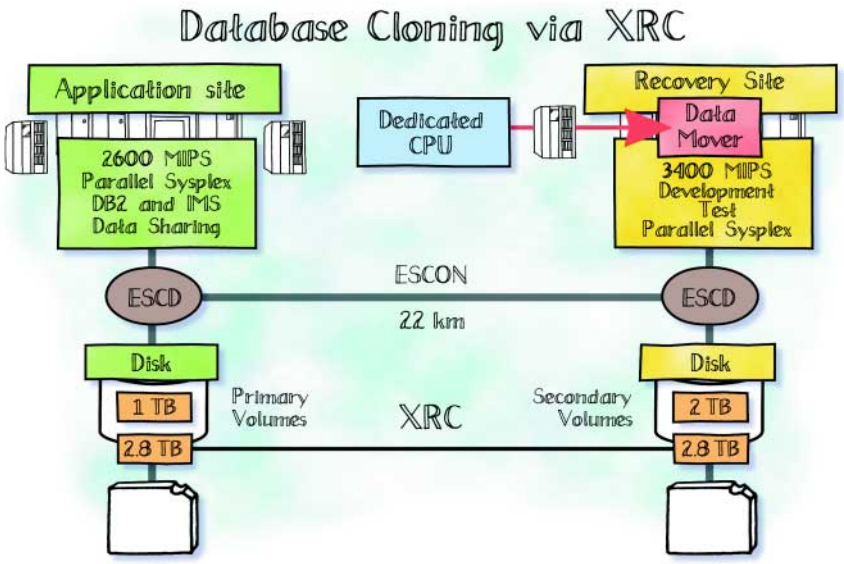

XRC is transparent to applications. XRC is controlled via TSO/E commands.

Who Will Need the Enterprise Storage Server?

The Enterprise Storage Server (ESS) is the storage of choice for those wanting the ultimate in enterprise disk-storage consolidation and data sharing on multiple or heterogeneous server platforms with combined storage requirements of 420 GB or more.

The ESS is the longer-term strategic enterprise disk-storage solution for those who want to manage all of their storage from a single location. ESS is the right choice for those requiring high-performance RAID-5, read and write cache, or the flexibility of common storage for multiple servers with ESCON, FICON, SCSI and Fibre Channel attachments.

The ESS supports rapid universal access to vast quantities of data through many advanced functions and features, making it a workhorse for support of business-intelligence and other business-critical applications. Because of its enterprise-wide support and management scope, ESS is tailor-made to help provide consistent, efficient and effective enterprise-resource planning. Using ESS to address any or all of your strategic and tactical business initiatives will give your organization the business advantage needed to survive and thrive in the e-world.

The Enterprise Storage Server is *the* server for the next millennium.

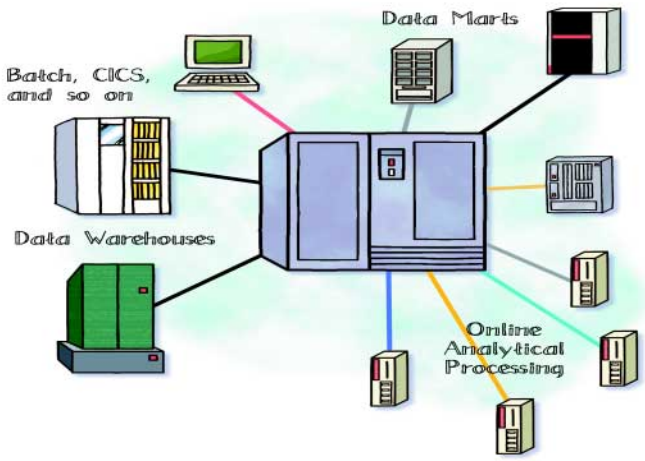

Kim Frank (kcfrank@us.ibm.com) has 28 years of in-depth experience in the Information Systems industry with positions in operations, programming, systems and programming, application development, marketing, sales, and management. He has designed, implemented and managed complex online application development projects. He has also managed technical professionals in marketing and support, resulting in 100% customer satisfaction, and managed 70 operations and technical support professionals. He has extensive technical design and implementation experience in large, complex environments with multiple processors, large networks, and terabytes of storage, with high availability requirements.

Enterprise Storage Server Fact Sheet

Overview

The introduction of the IBM Enterprise Storage Server (ESS) affirms IBM's leadership among disk-storage providers. This advanced, SAN-ready disk- storage system provides outstanding performance, scalability, and universal access across all major server platforms.

The ESS uses IBM's Seascape architecture with advanced hardware and software technologies to deliver breakthrough levels of performance and to maximize data sharing across the enterprise. The ESS provides the ultimate in scalability and flexibility across many platforms and configurations

Highlights

Extensive Connectivity

- Simultaneous attachment to a wide variety of servers, including UNIX, Windows NT and Windows 2000, Novell NetWare, AS/400, S/390 servers
- Fast data-transfer rates with a wide range of attachment technologies
- Fibre Channel, Ultra-SCSI, ESCON, FICON

Extreme Scalability

- 420 GB to 11.2 TB capacity
- Non-disruptive capacity upgrades
- Scalability for fast-growing environments (for example, e-business, Web-based commerce, business intelligence, transaction processing, enterprise resource planning)

Breakthrough Performance

The ESS combines state-of-the-art hardware design with software enhancements to provide leading-edge performance.

Performance-Optimized Hardware

- Large cache delivers high performance and helps protect data with mirrored write cache.
- SSA-160, the industry's leading non-arbitrated serial disk-attachment architecture.
- IBM 10000-rpm drives deliver rapid data access and fast data-transfer rates.
- Striped RAID-5 architecture optimizes performance for many different workload types.
- Two powerful, four-way SMP RISC processors.

Unique S/390 Performance Functions

Parallel Access Volumes, Multiple Allegiance, and I/O Priority Queuing practically eliminate contention, dramatically improving performance, match disk performance to workload priorities, and significantly reduce costs by simplifying storage management.

Comprehensive Availability

System, data and application availability anytime, anywhere.

System Availability

- Three-year warranty.
- The ESS uses redundant components.
- Dual active processing complexes provide backup to each other.
- RAID-5 disk arrays are designed to protect data against disk-drive failures.
- Redundant design also facilitates non-disruptive maintenance.

Data/Application Availability

- The ESS provides functions that aid in protecting and managing data on S/390, UNIX and Windows NT systems.
- Peer-to-Peer Remote Copy (PPRC) and Extended Remote Copy (XRC) provide protection against data loss by storing a remote second copy of data.
- FlashCopy copies data almost instantly for backup, testing, development or data mining.
- Storage partitioning allows storage capacity to be divided among host systems, as required.

Investment Protection

- Previous-generation IBM 7133 Serial Disk System capacity can be used with the ESS, extending the life of this investment.
- Seascape architecture allows easy "snap-in" additions and upgrades with new technologies.
- SAN-ready: The ESS is ready today for tomorrow's SAN environments.

Storage Management

IBM StorWatch family of software provides centralized Web-based management of Enterprise Storage Servers wherever they are located.

Specifications for Hardware Components

- Dual four-way SMP RISC processors.
- 8-GB cache/384-MB nonvolatile write cache
- Supports 9-GB, 18-GB, and 36-GB serial disk drives plus most IBM 7133-020 and 7133-D40 Serial Disk System drawers.
- From 4 to 32 server connections using Ultra-SCSI, ESCON, FICON, or Fibre Channel technology (or an intermix of these).
- Disk arrays can be used as RAID-5 or non-RAID (also known as *JBOD*).
- Unsurpassed availability.
- Redundant components: power, cooling, adapters, buses, processor complexes and disk drives.
- Reduced costs of maintenance.
- "Call home" capability places service call automatically when a problem arises.
- Concurrent maintenance enables maintenance while disk system is running.
- Concurrent microcode activation.
- Single-phase or 3-phase power with battery backup.

What's New with ACF/NCP and SSP

by Preston Johnston

NCP is 27 years old. It is a software product that had its genesis in the 70s when start-stop and bisync were all the rage, and it is still going strong in the new millennium—the age of the Internet. NCP is certainly one of the longest-lived software products of all time, making it a survivor in this ever-more-rapidly changing IT world.

To keep up with the ever-changing IT world, NCP has changed immeasurably over the years, evolving into one of the most reliable and trusted software products from IBM ever, being the cornerstone of SNA networks.

Even today, SNA networks are valued for stability and dependability, which companies continue to require to communicate highly valued information assets.

We are making a service update package (SUP) release available that includes a functional enhancement and the latest service updates for NCP V7R8 and SSP V4R8, originally shipped in August '99. A SUP release is available for OS/390® and VM users. The enhancements are also available as individual APARs for OS/390, VM and VSE, which will allow you to choose the most optimum manner of installing these enhancements.

Many of these new enhancements available for NCP and SSP are the direct results of customer requests. This is an another example of IBM's continuing commitment to SNA and to the NCP product family.

The following section provides an overview of the new functions. You will find a matrix at the end of this article that contains the APAR information for each new NCP/SSP function. This will facilate ordering individual APARs if you choose to install individual new functions rather than the entire SUP tape.

Token Ring Enhancements

Token-Ring Outbound Data Queue Length

NCP V7R8 extends performance monitoring used today to collect data queue lengths for subarea stations to now include token-ring stations attached to the 3746 Nways® Multiprotocol Controller Model 900. This enhancement, when used in conjunction with NTuneMON™, allows additional monitoring and tuning of network performance.

Predefined Token-Ring Route

NCP V7R8 provides you with a means to override the normal route-discovery process for a 3745 Communication Controller (TIC1/TIC2) token-ring subarea connection. At system definition time, you can specify the exact bridged route that must be used to establish the connection. When this NCP option is selected, the connection will not become operative until the specified route is available, thus ensuring that only the predetermined optimal route is used.

Token-Ring Congestion Control Enhancement

The DYNWIND token-ring support originally supported for 3745-attached token-ring stations is now available for the 3746-900. Dynamic windowing lets you specify how much to reduce a link station's frame queue before requiring an acknowledgment when frame loss is detected. By making this a variable, it allows the network to recovery more quickly from lost-frame conditions in both normal and congested environments.

Network Management Enhancements

Switched SDLC Subarea Disconnect Enhancement

NCP V7R8 will now report completion of an operator-initiated hang-up from the secondary side of a switched SDLC subarea connection with a VTAM® message. This allows automated host-operational procedures to be aware of the link status constantly and not just on an exception bases.

LUs Alert Threshold Enhancement

NCP V7R8 now separates dynamic LUs from static LUs, which allows you to define a common alert threshold for LU utilization for all NCPs in a network, regardless of differing LU pool sizes specified in different NCPs in the network. This enhancement simplifies network management by ensuring that any LU utilization alert received will indicate a common threshold and will allow for common recovery procedures for all NCPs in a network.

DYNPOOL Control Block Recovery

NCP V7R8 will check the validity of all DYNPOOL control blocks prior to freeing the buffer. With this enhanced first-failure data capture, service time can be shortened by ensuring that the buffer contents are not lost in case the control-block contents are invalid, something that indicates a possible problem.

Additional Diagnostic Codes

NCP V7R8 provides additional diagnostic codes to be added to protocol error messages for 3746-900 connections. This enhancement will assist in pinpointing the cause of any protocol error reported to the 3746-900.

NCP Performance Enhancements

Frame-Relay IP Congestion Control Enhancement

By expanding the amount of NCP buffers that can be used by a frame-relay station, IP throughput is improved. This will help to prevent one protocol locking out another, thus improving NCP's overall throughput.

Miscellaneous Enhancements

Expanded X.25 Network Support

NCP V7R8 now allows the optional definition of delayed- or immediate-packet acknowledgments. This will allow greater flexibility because some X.25 networks do not allow delayed-packet acknowledgments (piggybacking).

ALCI Extended Application Name

The application name (LOGAPPL) keyword will now allow network-qualified names up to 17 characters in length.

Preston Johnston (prestonj@us.ibm.com) joined IBM in 1969, retired in 1997 and returned part-time in 1998. Preston's first involvement with NCP was in 1971 with the testing of the initial releases of NCP. After many years' absence from NCP, he has returned and is now a part of the Business Line Management organization.

NCP/SSP APAR Ordering Information

Enhancement	NCP VM MVS™	NCP VSE	SSP MVS	SSP VM	SSP VSE
TR Outbound Data Queue Length	IR42166	IR42302			
Predefined TR Route	IR42305	IR42306	IR42372	IR42376	IR42379
DYNPOOL Control Block Recovery	IR42308	IR42321	IR42373	IR42378	IR42379
LUs Alert Threshold	IR42460	IR42467			
Switched SDLC Subarea Disconnect	IR41936	IR41941			
Dynamic Windowing for TIC3	IR42608	IR42611	IR42617	IR42618	IR42619
Partial Dump Format			IR41661	IR42664	IR42665
FR IP Congestion Control	IR42059	IR42060			
Diagnostic Codes	IR42370	IR42375			
Internal NCP Traces	IR42599	IR42601			
Expanded X.25 SupportIR42600	IR42602	IR42660	IR42662	IR42663	IR42663
ALCI Extended Appl Name	IR42545				
Extended SIT Trace Data	IR42546	IR42548			

Get More NTuneMON for your Money

by Dan Harker and Carl Marlowe

NTuneMON™ Version 3 is a significant tool for managing 3745 and 3746 communication controller networks. Prior to Version 3, NTuneMON Version 2 enabled you to view over 400 NCP generation parameters and, for an additional monthly fee (per controller), the optional NTuneNCP Tuning Feature enabled you to dynamically tune over 200 NCP generation parameters.

New with NTuneMON Version 3, the optional NTuneNCP™ Tuning Feature is now included in the base NTuneMON product. This new packaging now offers you both monitoring and tuning capabilities for an unlimited number of 3745 or 3746 controllers at one low monthly fee.

This offering provides you with the maximum ability to manage and dynamically tune NCP resources in many communication controllers while reducing your overall cost of ownership.

What's New with NTuneMON

NTuneMON V3R1 will supply additional support functions for the new items being added through program temporary fix (PTF) to NCP. NTuneMON will also provide additional functions in the areas of NCP service aids and IP-NCST connections.

The new NCST Status Information Detail panel includes information on the status of a session, such as an LU-SSCP session: whether the session is active or inactive, primary or secondary. It will also indicate whether a forced deactivation is in progress. In other words, it tells you what the status of the NCST-LU session is. This panel also provides information on the PLU-SLU session status, on the BIND, and on activation or deactivation of the PLU-SLU session.

Information is also available on the pacing status of the session, including the current pacing limit and count. This panel supplies the interface address and name, and you will also be able to see the remote session partner name, subarea value and log mode name, if those are available in the NCP.

The NCP has some internal traces, which are usually activated at the request of the NCP Level 2 Support group. Prior to NTuneMON V3R1, if you wanted to modify, activate or change these traces, you first had to take a dump of the NCP or the link edit map of the NCP gen to locate the address of the trace location in the module in which it resides, and then modify NCP storage from the MOSS console.

Now, with the NCP APAR and NTuneMON V3R1 installed, the new panel, ATUSX NCP Trace Functions, which you access by using a PF key from the ATUST NCP Function Activation panel, allows you to activate, change or modify the trace parameter for the following traces and service aids:

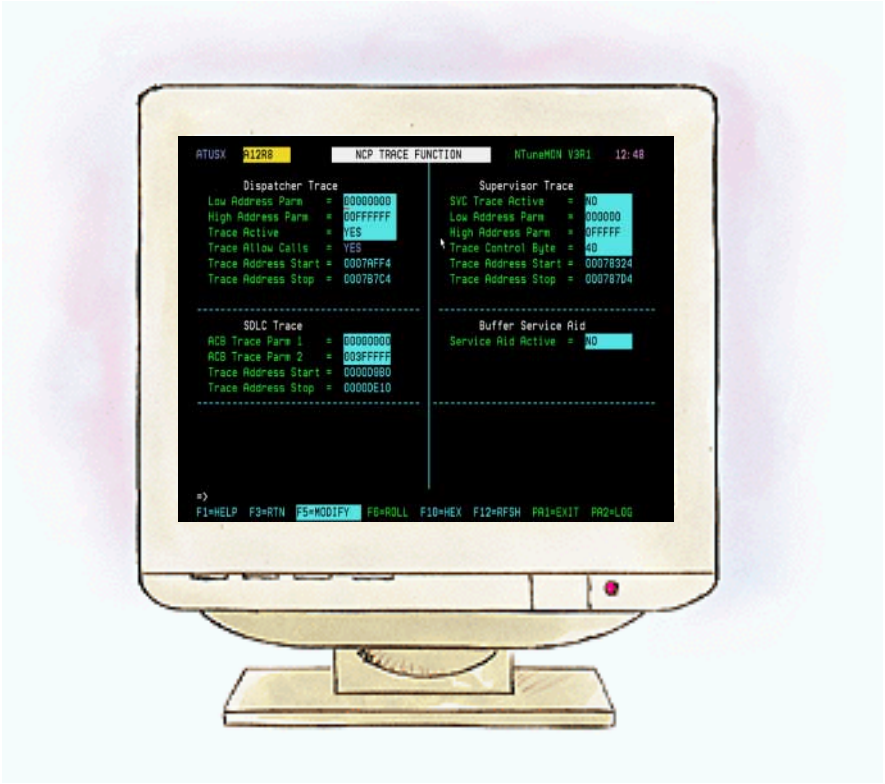

- The Dispatcher trace, which runs all the time the NCP is running. You can modify it to trace only selected addresses for diagnostic purposes in order to isolate problems with the NCP code. You do this by modifying the high and low addresses of the area to be traced.
- The SVC trace which, when active, traces the function of calls to the supervisor area of the NCP code, tracing a selected set of registers, or can be modified to trace all of the registers along with the address from where the SVC was called. You can modify the address range or restrict it by setting the high and low address range. You can use this trace to diagnoses potential NCP problems.
- The SDLC Level 3 I/O trace, which you can use to find a problem with a link connection or the NCP code that can affect the operation of that connection. You set the trace by storing the ACB address or addresses of the link to be traced. To simplify this process, we have modified the ATUPL panel to provide the ACB addresses whenever the PF function key is used.
- The Buffer Service Aid function, which helps you to isolate the cause of a buffer abend problem. This can either be a problem in the NCP code or a hardware problem when the hardware causes an overlay of data into an NCP buffer.

In response to customer requests, we have added an additional function to the ATUGP Network CB Pools/Tables panel: you can now modify the value for NETLIM, which controls the number of SNI session that can be set up across a network. This will temporarily allow sessions until a new NCP gen can be run.

We have also made a number of enhancements to the ATUTL Token-Ring SNA Station panel. To start with, we made 12 additional fields tunable for TIC3 adapters. These fields are: T1LOC, T1REM, T2LOC, T2REM, PIUs/BLOCK, MAXDATA, BLOCK GEN, MAXOUT GEN, N3GEN, and the DYNWIND parameters NW, DW and DWC.

Additionally, the ATUTL panel includes the new Outbound Queue Length and PreDefinedRoute parameters. The Outbound Queue Length field displays the number of elements currently queued to the logical link outbound queue for peripheral logical stations. If the PreDefinedRoute field is set to yes, the PF key is displayed. This key can bring up the ATUTL2 Token-Ring SNA Station Panel 2, which allows you to modify the Token Ring Bridged Routing information. This is useful where a specific bridged route is required between subarea nodes rather than the route determined by the token-ring route selection process.

On the ATUXP X.25 MCH Physical Line Details panel, we have added a new parameter for X.25 ODLC lines, PLPPIGGYB, which reflects the value as coded on the X.25 MCH statement.

Finally, on the ATUTP1 Token-Ring TIC Details Page 1 panel, we have made the DYNWIND DW and DWC tunable for TIC3 physical lines. This provides consistency with the NW parameter, which was already tunable from the ATUTP1 panel.

Dan Harker (dharker@us.ibm.com) joined IBM in Field Engineering in 1966. He has worked in the NCP Service organization since 1984 and on the NTuneMON Development Team since 1992.

Carl Marlowe (cmarlowe@us.ibm.com) joined IBM in 1980 and has worked as a development programmer on various products for the finance and manufacturing industries. He joined the NCP Design and Development Team in 1997 and is currently working on NTuneMON.

How We Test New Program Enhancements

by Sue DeMarrais and Jeanette Baugh

When NCP, ACF/SSP, and NTuneMON™ incorporate program enhancements, the test organization gears up for the task of validating the updated products. This includes not only verifying new individual functions, but also regressing existing functions, testing the packaging of the total release, and confirming that APARs are processed correctly. We test the enhancements in a variety of environments, and in conjunction with numerous networking products.

Functional Verification Test

The first stage of testing—Functional Verification Test (FVT)—concentrates on the newly introduced enhancements. At this stage, we devise specific test-case scenarios and then review them for accuracy and completeness. Once they are approved, we produce testing environments with NCP generations that are created using the updated ACF/SSP and NCP products. We validate ACF/SSP keyword additions and changes using maximum, minimum, default, and numerous representative values. To prove that new functions work as intended, we then exercise NCP and NTuneMON enhancements. After completion of the FVT, the test-case scenario documents are updated with the results, giving future test efforts a repeatable process.

System Verification Test

The second stage of testing—System Verification Test (SVT)—brings the updated product into a larger, more complex network environment. At this stage, we utilize the NCP product set in a multiple-host network with an assortment of applications and protocols. In addition to implementing and exercising new functions, we generate a variety of traffic types in the network, including pre-SNA, SNA, APPN®, HPR, and IP, and process them through the NCPs. We vary data sizes and traffic loads, and generate thousands of sessions with real and simulated devices. We test stress, error, and recovery processing concurrently with high traffic levels in scenarios such as:

- Unplugging and replugging cables, adapters, couplers
- Powering devices off and on
- Causing beaconing on token-ring interfaces
- Deactivating and reactivating lines, PUs, and LUs repeatedly
- Shutting down user applications during session processing
- Forcing termination of sessions
- Causing automatic network shutdown (ANS)
- IMLing the attached 3746-900 during operation
- Performing concurrent maintenance on 3746-900 resources
- Causing slowdown processing
- Causing path switching of HPR sessions

The SVT integrates many software products and protocols, including:

- ECHO
- eNetwork Personal Communications/ Communications Manager
- FTP
- Linux
- Microsoft® Windows NT® and Windows® 98
- NCPRoute
- NPM
- Operating System/2® (OS/2®)
- OS/390®
- RouteD
- SNALINK
- TCP/IP
- Teleprocessing Network Simulator (TPNS)
- Tivoli® NetView®
- TN3270
- Assorted workstation products

It also integrates various hardware products, including:

3745

3746-900

2210

3725

3174

3270

3720

2740

3101

PCs

S/390®

Simulated TPNS devices

At the end of the SVT, all network definitions are saved for subsequent test efforts.

Additional Testing

For a new release, the Packaging & Release (P&R) Test follows the SVT. At this point, we collect all code changes for the release, build them into a new installation tape and install the tape in new libraries. We then regenerate the NCPs from the SVT against these P&R libraries. In addition, we use the updated NTuneMON libraries. We use the same network configuration implemented for the SVT to authenticate that the product has been built correctly. The P&R covers many of the same scenarios as the SVT, to ascertain that the newly packaged product is equivalent to the previously verified one.

As a supplement to an existing release, enhancements may also become available as individual APARs. In this case, we carry out a PTF Test to verify the APAR processing. PTFs are applied over a base level of code, in an order determined by necessary prerequisites. When a PTF is processed, the tester reviews the cover letter, checks to see if a specific test is requested, then receives and applies the modification. In the MVS™ environment, this application process is automated through System Modification Program/Extended (SMP/E). VM PTFs are applied through VM/Serviceability Enhancement Stage (VM/SES) or manually, depending on the release. VSE PTFs are applied with Maintain System History Program (MHSP). Once the PTF applies successfully, we replace NCPs from the SVT configuration with NCPs generated against the PTF libraries and run a regression test. When we have verified a PTF and ensured that any special test requirements are met, the APAR processing is completed and the PTF is made available to our customers.

Where the Enhancements Are First Used

After the new NCP, ACF/SSP, and NTuneMON products become available, one of the first environments to utilize them is the IBM System Test Lab, Raleigh, North Carolina. S/390, TCP/IP, and numerous workstation test groups directly utilize new NCP, ACF/SSP and NTuneMON features.

NCPs are also used as a production backbone for many workstation and e-business test groups, working with products such as Communications Manager, Host On-Demand, and Communications Server for Windows NT. These groups use NCP as a workhorse, to provide around-the-clock availability and performance. The Raleigh System Test Lab is continually improving its test processes with better test-tracking tools, shared network configurations, and improved traffic-generation tools.

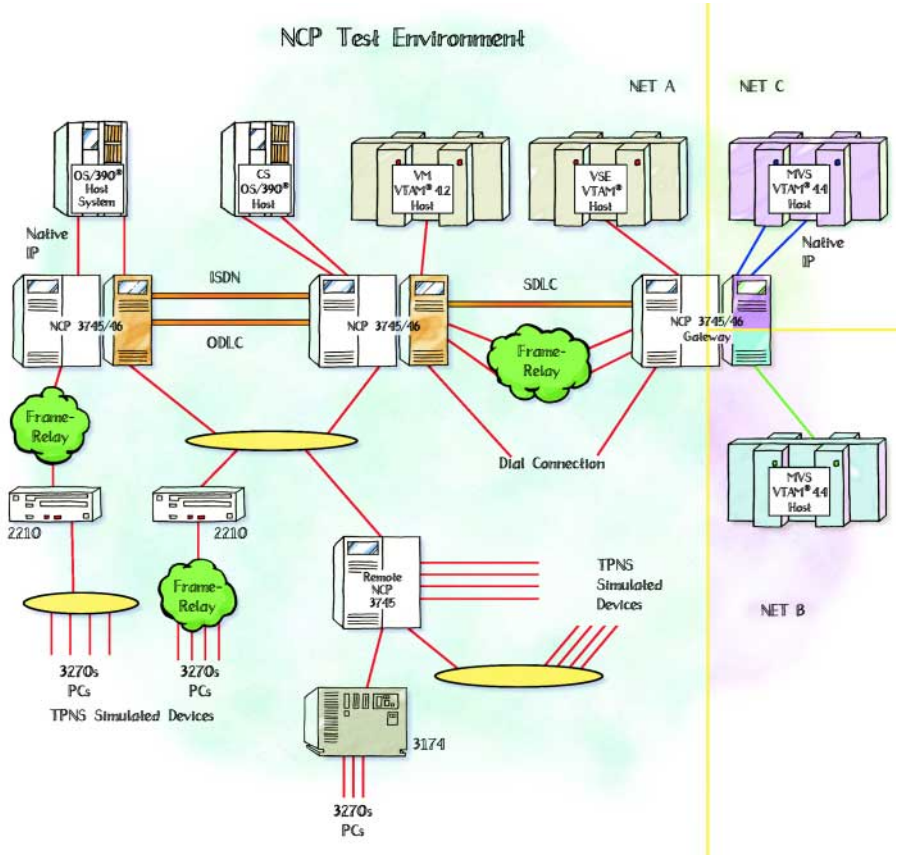

Sue DeMarrais (demar@us.ibm.com) joined IBM as a System Verification Tester for the NCP product set in 1990. She has participated in testing several NCP, 3746 Model 900, and S/390 releases.

Jeanette Baugh (baugh@us.ibm.com) joined Office Products Division in 1976 in Detroit, Michigan. She relocated to Portland, Oregon in 1979. In 1989, she transferred to Research Triangle Park, North Carolina, as a Level III for NPM and an SVT tester for NCP. Since 1995 she has been a PTF tester for NCP, VTAM and OS/390.

3746 Nways Multiprotocol Controller Year 2000 Enhancement Overview

by John Foley, Jr.

As 1999 came to an end and the clocks turned to January 1, 2000, it was our number one goal to help you transition your networks into the next millennium with little or no disruption to normal business operations. Now that it's over, it's time to quickly turn our eyes toward the future and ask three questions:

1. "What's the future for SNA networks?"
2. "How can I implement a network design that allows me to smoothly integrate my mission-critical SNA applications with Web-based Internet applications over a single IP backbone?"
3. "Can I maximize my current 3745/46 hardware investments and use these footprints to help bridge myself to the next major step in network computing?"

The answers to the first two questions are explored throughout articles within this magazine.

The answer to the third question is a resounding *yes*!

Year 2000 has truly been a busy year for the 3746 team. In the first 6 months, we've traveled the world, reinforcing the fact that IBM is committed to delivering SNA solutions that help you successfully evolve some of the world's largest SNA networks into "Lean and Mean e-business trafficking machines."

Recognizing that gigabit switching technologies are upon us, the 3745/46 and NCP development team were compelled to implement enhancements that provide you with additional performance, capacity and management functions. It is our goal that these enhancements provide the functions that allow you to fully capitalize on your current 37xx investment. Incorporating these updates provides you with the time necessary to implement a strategic plan that evolves your network-of-today into the network-of-tomorrow, providing both Internet and intranet solutions that maintain and implement the attributes and benefits of both SNA and IP environments.

On February 8, 2000, IBM announced two new processors for the 3746 Nways® Multiprotocol Controller. Here is an overview of the enhancements:

- A new Service Processor Type 4 (SP4) replaced the Service Processor Type 3 for the IBM 3745 Models 17A, 21A, 31A, 41A, 61A, and 3746 Models 900 and 950. The new processor includes a control unit equipped with a Pentium® III 533-MHz processor, 128-MB memory, 13.5-GB hard-disk capacity, a CD-ROM drive, a 15-in. monitor, and an IBM 7858 Modem for remote operator access at speeds up to 56 Kbps. This new processor enables fast operator intervention, with a response time reduced by a factor of up to five, compared with older versions of the Service Processor.
- A new Network Node Processor Type 4 (NNP4) replaced the NNP Type 3 for the IBM 3746 Models 900 and 950. The new processor includes a Pentium® III 533-MHz processor, 128-MB base memory, and a 13.5-GB hard-disk capacity, which provides faster network restarts. In addition, the NNP Type 4 memory can be expanded from 128 MB to 256 MB.

In addition, we are offering the following capacity and performance enhancements for the 3746 Nways Multiprotocol Controller, which were announced in June 2000 and are planned for availability starting the second half of 2000:

- *50% more system services control point-to-logical unit (SSCP-LU) sessions on NNP3 and NNP4*

  The SSCP-LU control session capacity is increased by 50%, from 40000 to 60000 sessions. This allows migration from NCPs with a large number of defined LUs.

- *33% more logical unit-to-logical unit (LU-LU) data sessions on NNP3 and NNP4*

  The total number of LU-LU sessions (APPN® or DLUR) routed by the 3746 Network Node is increased from 30000 to 35000. Up to 20000 of these LU-LU sessions can be established by the NNP, an increase of 33% compared with the previous maximum of 15000 sessions. The other LU-LU sessions (up to 35000) are intermediate sessions established by other network nodes.

- *33% more physical units (PUs) on Token-Ring Processor Type 3 (TRP3)*

  TRP3 can now connect 4000 physical units (PUs). Previously, the maximum was 3000 PUs. The TRP Type 2 can connect up to 2000 PUs. The increased TRP3 connectivity applies to PUs activated by NCP, and PUs activated by the NNP, in any mix up to 4000. This enables an increase in the number of end users supported on existing token-ring processors.

- *33% more IP routes on Controller Bus and Service Processor Type 3 (CBSP3), increasing the maximum number of IP routes from 7500 to 10000.*

- *33% more Open Shortest Path First (OSPF) external routes on CBSP3, increasing the maximum number of OSPF external routes from 7500 to 10000.*

- *Increased IP throughput on ESCON® Channel Processor Type 3*

  The maximum transmission unit (MTU) size for IP traffic over ESCON channels is increased from 4 KB to 8 KB, to enable higher channel throughputs for large-volume data transfers. This allows faster file transfers, ESCON-to-ESCON, or token-ring LAN-to-ESCON. The actual performance information will be available June 30, 2000, at: **ibm.com**/networking/3746

- *Voice traffic over frame-relay lines connected to Communication Line Processor Type 3*

  Frame-relay lines carrying voice traffic (Voice over IP or Voice over Frame Relay) or voice and data traffic (SNA, IP, and so on), can be connected to the CLP3. This support, which is available for lines activated by either NCP or the NNP, enables voice and data consolidation over an existing network infrastructure. Frame-relay connections and SDLC leased lines (redefined as frame-relay lines) can be used between 3746 sites to save the costs of separate leased lines dedicated to voice traffic.

- *Dynamic Windowing on Token-Ring Processor Type 3*

  The new algorithm allows the NNP and NCP to regulate traffic more efficiently over token-ring ports.
  The benefits of this mechanism are:

  ✓ More effective utilization of WAN network bandwidth after detection of network congestion

  ✓ Faster recovery from network problems after detection of frame loss

- *Controller Configuration and Management (CCM) Enhancements*

  ✓ Dynamic discovery of low-entry networking (LEN) nodes and network nodes (NNs) without control sessions
  These nodes are automatically discovered without any PU definition in CCM. The node-initiated sessions do not require any LU definition in CCM. Only the VTAM® software-initiated sessions require the corresponding node LUs to be defined in CCM.

  ✓ Dynamic change of maximum number of PUs per TIC3 port

  Up to now, the operator had to stop the port to change the maximum number of incoming calls per TIC3 port. With this improvement, the users already connected are not affected by a change in the maximum number of PUs per TIC3 port. In a multi-3746 installation, this allows the operator to non-disruptively adjust the total number of PUs per 3746/NNP, depending on the number of operational 3746/NNPs. As an example, a data center with two 3746s and 5000 PUs would be defined in such a way that a maximum of 2500 PUs is accepted per 3746. In normal operations, this allows the connections to be evenly spread between the two 3746s, even if the 3746s started one after the other. If one 3746 is not available, the operator can dynamically change the TIC3 definition of the active 3746 to increase the total number of PUs to 5000.

  ✓ Editing of CCM configuration files (CCM "batch")

  For APPN/HPR/DLUR definitions, this method allows very easy replication of parameter definitions and it can be much faster than the standard CCM graphical interface. When you use Distributed Control Access Facility (DCAF) on a workstation running the OS/2® operating system, you can download CCM configuration files from the Service Processor of the 3746, and then extract the APPN configuration files and convert them into ASCII files. You can use a text editor, either locally on the OS/2 workstation or on another workstation (Windows®, OS/2, and so on), to complete or update the APPN ASCII configuration files. The OS/2 workstation is finally used to convert the ASCII files back to CCM file format and then to upload them to the Service Processor via DCAF.

- *Network Management from Tivoli® NetView® for OS/390®*

  The following RUNCMDs are added or enhanced for 3746 control using NetView software:

- Activate/deactivate IP resources
- Display the 3746/NNP counters of PUs and sessions (LU-LU, SSCP-LU)
- List all ports sorted by name
- List all stations sorted by name
- List a summary of the APPN sessions per alias name
- List a summary of the topology per APPN network node
- List all the topology data for a given APPN network node
- List a summary of the directory per APPN network node
- List all the directory data for a given APPN network node
- Request the APPN control program dump
- Monitor the APPN control program dump request completion

- *Operations Management*

  ✓ Four concurrent operators using Java™ consoles

  Up to four operators can have concurrent access to the 3746 through Java consoles. The first one is provided with a full-access privilege. The other three can have either a full-access privilege or a view-only access privilege. With this multiple-operator access to the 3746, a help desk can rapidly respond to end-user requests.

  ✓The user logon for Java-console access over switched networks is now customizable, which allows 3746 users to choose simple or existing passwords.

  ✓Flexible password definition for Telnet access and CCM

  The Telnet access and CCM password is no longer restricted to a fixed length of 8 characters. This allows 3746 users to extend existing passwords to 3746 CCM/Telnet and to have a single password for all routers in the network.

All of the new functions listed in the June 2000 announcement are available as a new "3746 Extended Functions 5" Licensed Internal Code feature of the 3746. For more details and prerequisite information, refer to the February and June 2000 product announcements at:
ibm.com/networking/3746

John Foley, Jr. (foleyjoh@us.ibm.com) has been a member of the NCP organization for 10 years. His responsibilities include world-wide marketing & sales of the 3745, 3746, NCP and its associated family of products.

NTA: A Service Tool for Analyzing Complex Performance Problems in TCP/IP and SNA

by Bob Springsteen

The network is crucial to communications for any large corporation. The importance of availability and performance is clear. Long response times or slow data throughput cause headaches for network support personnel. Today's e-business network is primarily TCP/IP, although enterprises with SNA subarea implementations are also recognizing performance impacts, especially those migrating to the APPN® environment. This article discusses performance challenges for these environments and offers a successful method for analysis and resolution based on end-to-end and session-level performance analysis.

Identifying the problems

To quickly identify where incorrect settings are causing users to experience poor performance, begin with an end-to-end analysis. Analyzing the behavior of the network requires the use of tools that provide a statistical representation of operational characteristics to identify problems, evaluate network performance and guide actions for tuning. IBM Performance Management and Capacity Planning Services provides a network traffic analysis (NTA) tool and consultative support targeted at improving TCP/IP and SNA network performance. NTA uses trace data captured from system trace tools as well as standard network data-capture tools. Analysis is automated for subarea data flows as well as TCP/IP packet traffic. A view of end-to-end routes experiencing problems provides a starting point for the performance analysis.

Effective analysis of the TCP/IP environment requires an understanding of the fundamental physics of network behavior, control mechanisms for TCP/IP performance, and identification of performance problems through network-behavior analysis. This includes an examination of TCP flows, packet sizes, window behavior, flow and congestion controls, and data-link operations.

A high-level view of all the routes in the flow is the best place to start your analysis and helps to identify those routes that are experiencing problems. From this perspective, you can identify and focus on the cause of performance problems. The cause might simply be a need to make logical parameter changes or a routing adjustment to distribute the load and avoid overloaded resources. In some cases, the problem might be errors in a node or on a link. Once NTA has identified the problem, it can quicken resolution by automating the detailed analysis and providing specific recommended actions.

When you experience a performance problem, begin the analysis by tracing the segment of the network where the problem presents itself. The next step is to use the TCP/IP header data from this trace to understand the condition of the network at the time of the problem. You will need to examine both the connections that experienced problems and a packet-by-packet view of activity. Using NTA, you can identify those connections experiencing excessive response times and determine the cause of throughput problems.

Categories of Analysis

Analysis is divided into two categories: *global analysis and detailed analysis.*

Global Analysis

Global analysis is a high-level analysis of the trace data as a single entity. Results show a breakdown of the Layer 4 protocols (such as TCP and UDP) and the amount of data transported by each. The figure shows the division of the Layer 4 protocol as detailed by NTA.

Viewing the count of data packets by protocol helps you to verify that the trace data captured the problem you are experiencing. If the data does not match your expectations, you will need to take another data capture, perhaps in another area of the network, to complete a valuable analysis. Once you verify that the trace data includes network traffic at the time of the problem, you can continue your analysis by reviewing response times and other global information such as:

- Packet distribution by packet size with the packet count
- Trace statistic summaries
- Buffer-size distribution, including the receiver's window size with the packet count
- Round-trip time distribution with packet round-trip times including packet count The figure shows an example of response-time distribution.

DATA SUMMARY BY PROTOCOL

DATE	FILE	START	STOP	INTERVAL(in sec)
2000-04-25	demo	20:02:36	20:08:58	1

Display Problem Virtual Circuits | Display All Virtual Circuits

	Protocol Type	Count of IP Flows	Problem Flows	Byte Count of User Data	Peak Rate Bytes/Sec	Percent of Data
○	TCP	4	4	401823758	3816432	99.99
○	UDP	0	0	0	0	0.00
○	ICMP	0	0	0	0	0.00

Miscellaneous Protocols

Protocol	Number of Bytes	Number of Packets
89-OSPFIGP	43460	249

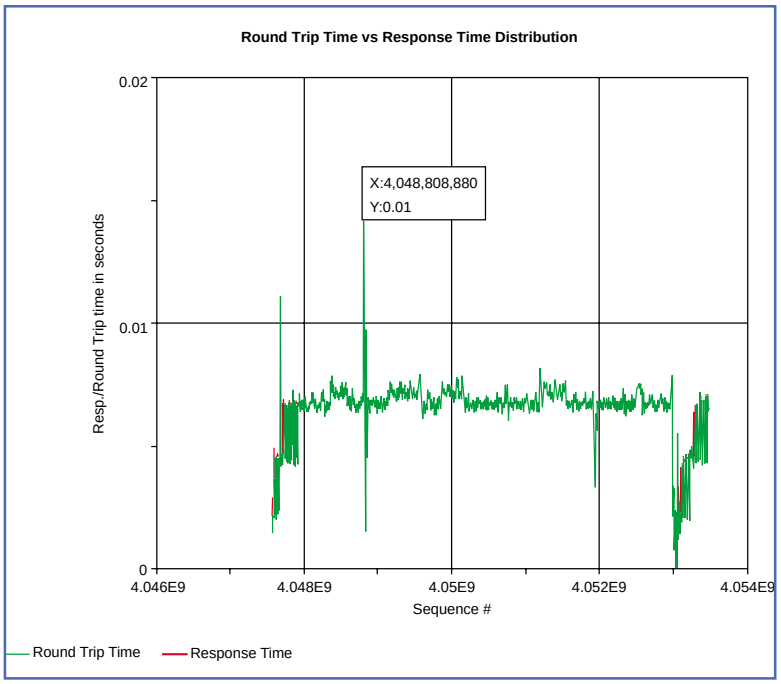

Detailed Analysis

Now that you have a high-level understanding of the network traffic conditions, detailed analysis begins. The detailed analysis provides an in-depth look at the data and helps identify problem connections. The following NTA results show the problem routes requiring more focused analysis. The severity of the condition of the route is highlighted by the color-coded circles, indicating the troubled connections.

LIST OF TCP ROUTES WITH POSSIBLE PROBLEMS

4 of 4 TCP CIRCUITS HAVE POTENTIAL PROBLEMS

DATE	FILE	START	STOP	INTERVAL(in sec)
2000-04-25	demo	20:02:36	20:08:58	1

Stat Report

	Ntwk	End Station	Perf	IP Source Addr:port# --> IP Destination Addr:port#	Dgrms	User Bytes	Long RTT	Long Resp	ReXmit	Dupl Rcds	Frag Dgrms	Re-Seq	Zero Window	Slow Rcvr	Reset Flag
○				010.228.064.082:00020-->153.002.017.006:01069	138700	0	0	0	0	6	0	0	0	0	0
○				010.228.064.082:00021-->153.002.017.006:01068	13	285	1	0	0	0	0	0	0	0	0
○				153.002.017.006:01068-->010.228.064.082:00021	11	139	1	1	0	0	0	0	0	0	0
◉				153.002.017.006:01069-->010.228.064.082:00020	277781	401823334	0	0	51	0	0	0	4	1	0

Stat Report

Further analysis of one of these probable problem routes provides specific statistical information for that connection. This example shows the detailed statistics for a probable problem route.

STATISTICAL DATA BY CONNECTION

153.002.017.006:01069-->010.228.064.082:00020

DATE	FILE	START	STOP	INTERVAL(in sec)
2000-04-25	demo	20:02:36	20:08:58	1

Display Data Throughput | Display Packet Trace | Clear All

Buffer Size Line Graph | Response Time Line Graph | Round Trip Time Line Graph | Response vs Round Trip Time

All Charts

	Time	Dgrms	Usr Dgrms	Usr Bytes	SegHdr Size	MinSeg Size	MaxSeg Size	AvrSeg Size	Min RTT	Max RTT	Avg RTT	Min RESP	Max RESP	Avg RESP	Re-Xmits	Dup Records	Frag Dgms	Re-Seq	Min TTL	Max TTL	Min UnACK	Max UnACK
☐	20:06:06	1685	1684	2438432	24	1448	1448	1448	.001480	.014266	.006691	.001480	.014266	.006747	0	0	0	0	64	64	1448	26064
☐	20:06:07	2637	2637	3816432	20	496	1448	1447	.000063	.008176	.006441	.001351	.008176	.006493	24	0	0	0	64	64	1448	26064
☐	20:06:08	2596	2596	3759008	20	1448	1448	1448	.001537	.016156	.006842	.001537	.016156	.006842	0	0	0	0	64	64	15928	26064
☐	20:06:09	2605	2605	3717016	20	212	1448	1426	.001806	.014147	.006780	.002160	.016231	.007020	0	0	0	0	64	64	4132	26064
☐	20:06:10	2676	2676	3719700	20	212	1448	1390	.003355	.019380	.006553	.003585	.019380	.007246	0	0	0	0	64	64	17376	26064
☐	20:06:11	2568	2568	3718464	20	1448	1448	1448	.004503	.014133	.006942	.004503	.014133	.006942	0	0	0	0	64	64	18824	26064
☐	20:06:12	2530	2530	3663440	20	1448	1448	1448	.001742	.016352	.007036	.001742	.016352	.007036	0	0	0	0	64	64	10136	26064
☐	20:06:13	2592	2592	3753216	20	1448	1448	1448	.004521	.014045	.006872	.004521	.014045	.006872	0	0	0	0	64	64	13032	26064
☐	20:06:14	2574	2574	3727152	20	1448	1448	1448	.005241	.008656	.006925	.005241	.008656	.006925	0	0	0	0	64	64	18824	26064
☐	20:06:15	2528	2528	3660544	20	1448	1448	1448	.000005	.011858	.003347	.000005	.011858	.003347	0	0	0	0	64	64	13032	34752
☐	20:06:16	2568	2568	3718464	20	1448	1448	1448	.000072	.006270	.001478	.000072	.006270	.001478	0	0	0	0	64	64	24616	37648
☐	20:06:17	2562	2562	3709776	20	1448	1448	1448	.000077	.009473	.001484	.000077	.009473	.001484	0	0	0	0	64	64	24616	37648
☐	20:06:18	2570	2570	3721360	20	1448	1448	1448	.000079	.003420	.001492	.000079	.003420	.001492	0	0	0	0	64	64	24616	31856

Further detailed information about specific connections that aid your analysis include:

- The number of times the receiver window size dropped to 0
- The longest response time
- Minimum segment size
- Number of fragmented datagrams
- Number of retransmits
- The maximum number of unacknowledged bytes outstanding

A packet-by-packet trace might also be required to analyzed data-flow details by connection.

Example: Poor performance and slow downloads

Performance challenges can result from incorrect system- and network-option settings. In an example TCP/IP environment, users were complaining of overall poor performance and slow downloads after a new version of the OS/390® operating system was installed. A trace of the problem showed that the TCP receiver's advertised window size was too large and retransmissions were occurring. The trace clearly showed that the OS/390 operating system's "window scale" option was in use. This gave the sender, the OS/390 operating system, permission to send enormous amounts of data into the network. The first router in the path did not have the buffer resources required to handle the data. Therefore, the router had buffer overrun, and packets were discarded. The sender would eventually resend the discarded packets, but the extra time involved severely impacted network performance. The solution was easy—simply turn off the default "window scale" option in the OS/390 software.

Traditional SNA Networks

Many organizations rely on well-performing SNA implementations to meet the communication needs of their business. When you begin the analysis of SNA subarea or APPN environments, you need to examine LU flows, packet sizes, window behavior, congestion, bottlenecks and data link operations. Transaction times in a network consist of processing times (service time) and delays. Some delays, such as propagation, cannot be controlled, while other delays are variable and controllable. You can focus on variable and controllable delays, especially your network queues, and know how to control them. Delays have a direct bearing on transaction times, which can affect expenses and user satisfaction.

One common performance problem in SNA environments is caused by blocked virtual routes (VRs). Usually, this is the result of data backing up in the network and therefore taking longer to get through the network. This blockage is referred to as subarea congestion and can be caused by a problem with the size of the VR pool or the VR pacing window sizes.

Analyzing VTAM® Internal Trace (VIT) data helps you identify the held routes. NTA automatically deciphers the trace data to highlight the routes that are experiencing poor performance. Drilling further down, NTA identifies potential causes of your problem along with recommendations.

Effectively identifying and resolving complex network performance problems in TCP/IP, SNA subarea, and SNA/APPN environments requires the ability to analyze network traffic quickly. This ability in turn requires skills, specific tools, and an efficient process. As the network continues to play an increasingly important role in successful business communication and operations, your ability to satisfy your users and your clients will be determined by your ability to identify and resolve performance problems quickly.

For more information

For more information on IBM Performance Management and Capacity Planning Services, call 1 800 426-4682 (in the U.S.A.), or 919 301-4141 (from outside the U.S.A.), or e-mail us at capacity@us.ibm.com. Visit **ibm.com**/services/tsm/nta to experience the capabilities of NTA.

Bob Springsteen (springs@us.ibm.com) works in the IBM Network Traffic Analysis (NTA) group. He is one of the original developers of the group. This group supports NTA clients and helps resolve network performance problems in SNA and TCP/IP using NTA's expert system analysis tools.

NTA support can be reached at 1 800 876-8801. The NTA Web page can be found at:
ibm.com/services/tsm/nta

Library Update

by Leyland King

A number of the NCP/SSP, EP and NTuneMON™ publications have been updated and made available in both hardcopy and softcopy for the NCP Version 7 Release 8 Supplementary Update and NTuneMON Version 3 Release 1. These include:

NTuneMON User's Guide	SC31-6266
NCP, SSP, and EP Resource Definition Guide	SC31-6223
NCP, SSP, and EP Resource Definition Reference	SC31-6224
LPS for NTuneMon V2 R6	CG31-6267

Along with these updates, the ACF/NCP, ACF/SSP, EP, NTuneMON and NPSI Softcopy Library Collection Kit, LK2T-0414-07, contains all other NCP library publications on CD.

Besides the BookManager® format, NCP V7R8 publications included with the supplementary update appear in PDF format on our CD as well as on the Internet. These publications are accessible using your Web browser. A pamphlet in the collection kit's CD jewel case tells you just how to access the information.

You can access these and all other available NCP publications on the Internet at: **ibm.com**/networking/ncp

Program directories for the supplementary update of the current releases of NCP and SSP as well as NTuneMON V3R1 are also available in softcopy on the NCP publications library CD as well as on the Internet in PDF format. You can also obtain them in hardcopy from Mechanicsburg as follows:

Program Directory for ACF/NCP V7R8 for MVS	GI10-6605
Program Directory for ACF/NCP V7R8 for VM	GI10-6607
Program Directory for ACF/SSP V47R8 for MVS	GI10-6606
Program Directory for ACF/SSP V47R8 for VM	GI10-6608
Program Directory for NTuneMON V3R1 for MVS	GI10-6609
Program Directory for NTuneMON V3R1 for VM	GI10-6610

NTuneMON Demonstration SW CD

A CD containing 90-day NTuneMON SW demonstration software is sleeved in the back cover of this magazine. Follow the directions on the sleeve to load the CD on your PC. You can order additional copies of the CD with copies of this magazine. See the back cover of this magazine for its order number.

As always, we invite you to contact us by e-mail through the online version of this publication on the Internet at: **ibm.com**/networking/ncp

Leyland King (leeking@us.ibm.com) joined IBM as an Information Developer in 1983. He has since participated in various product test, competitive analysis and information development efforts. As a software engineer and team leader for NCP publications development, he currently authors NCP library publications. He also manages the production and distribution of NCP and 3745/46 Today

User Group Meeting

Events in 2000	Date	City
AS/400 Technical Conference	June 12-16	Las Vegas, NV
CICS Technical Conference	June 12-16	Las Vegas, NV
Translation & Messaging Conference	June 12-16	Las Vegas, NV
MQSeries Technical Conference	June 12-16	Las Vegas, NV
VM/ESA and VSE/ESA Technical Conference In partnership with GUIDE SHARE EUROPE	June 26-28	Brussels, Belgium
Internet World	July 11-13	Chicago, IL
ISPCON Australia	Aug. 14-16	Melbourne, Australia
Solutions, The IBM Technical Developer Conference	Aug. 14-17	Las Vegas, NV
Solutions, The IBM Technical developer Conference	Aug. 14-17	Las Vegas, NV
Internet World	Aug. 15-17	Buenos Aires, Argentina
International Storage Symposium	Aug. 28-Sep. 1	San Diego, CA
ISPCON	Sep.	Frankfurt, Germany
Internet World	Sep.	Oslo, Norway
AS/400 Advanced Solutions—in partnership with COMMON EUROPE	Sep. 4-8	Strasbourg, France
Internet World at BITEC	Sep. 7-9	Bangkok, Thailand
Networking Solutions Technical Conference	Sep. 11-15	New Orleans, LA
RS/6000 Technical University	Sep. 11-15	Dallas, TX
Internet World	Sep. 20-23	Mexico City, Mexico
Internet World	Sep. 27-29	New Delhi, India
GSE Executive Training	Oct. 2-3	Palisades, NY
Secure World: IBM's End-to-End Security Conference	Oct. 2-6	Orlando, FL
Transaction & Messaging Congress—in partnership with COMMON EUROPE	Oct. 2-6	Amsterdam, Netherlands
Internet World	Oct 4-8	Dublin, Ireland
IBM UNIX Solutions Technical Update—in partnership with COMMON EUROPE	Oct. 9-12	Disneyland Paris, France
GSE Technical Symposium	Oct. 16-17	Montpellier, France
DB2 & Business Intelligence Technical Conference	Oct. 16-17	Las Vegas, NV
Events in 2000	Date	City
Internet World	Oct. 19-23	Milan, Italy
Common	Oct. 22-27	Baltimore, MD
MS Technical Conference	Oct. 23-26	Anaheim, CA
Fall Internet World 2000	Oct. 23-27	New York, NY
OS/390 EXPO & Performance Conference	Oct. 23-27	Washington, DC
Internet World Asia@Hong Kong	Nov. 1-3	Hong Kong, China
AS/400 Technical Conference	Nov. 6-10	Orlando, FL
Internet World	Nov. 7-9	Stockholm, Sweden
SPCON	Nov. 8-10	San Jose, CA
Internet World	Nov. 28-10	Paris, France
Internet World	Dec. 14-16	Tokyo, Japan
Internet World	Dec.	Caracas, Venezuela

Events in 2001	Date	City
Common	May 13-18	New Orleans, LA
Common	Oct. 21-26	Minneapolis

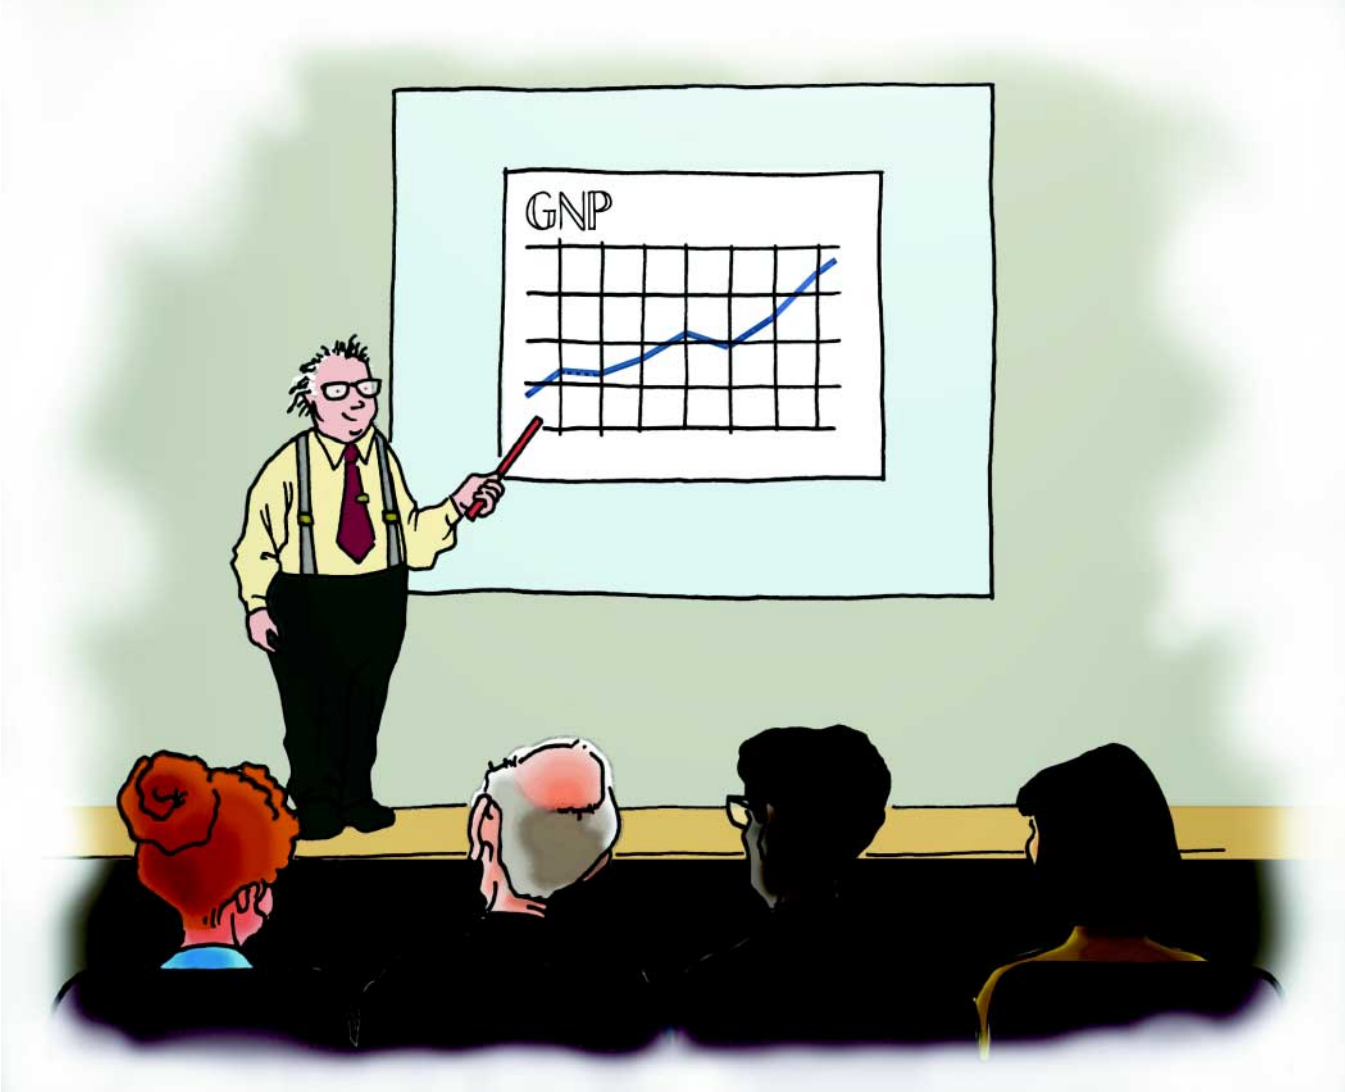

Post-Sales Product Support— What's New?

by Walt Reston

SSD Networking Business Unit: New Name, Same Commitment

We've changed our name, but not our commitment to provide you, our customers, with solutions, products and services that meet or exceed your expectations. The IBM System Storage Division Networking Business Unit (NBU), formerly known as the IBM Networking Hardware Division (NHD) continues to focus on providing you with high-quality SNA subarea, APPN® networking and token-ring products, *and world-class product support*.

Representative Networking Hardware products include:

NCP, Nways® Manager, 2210, 2212, 2216, 2220, 3172, 3174, 3745, 3746, 8230, 8239, 8270, 8271, 8272, 8274, 8275, 8371, 9729, and token-ring adapters.

We fully recognize that support, both prior to and after the sale, is a key element in your buying decisions—today *and* tomorrow.

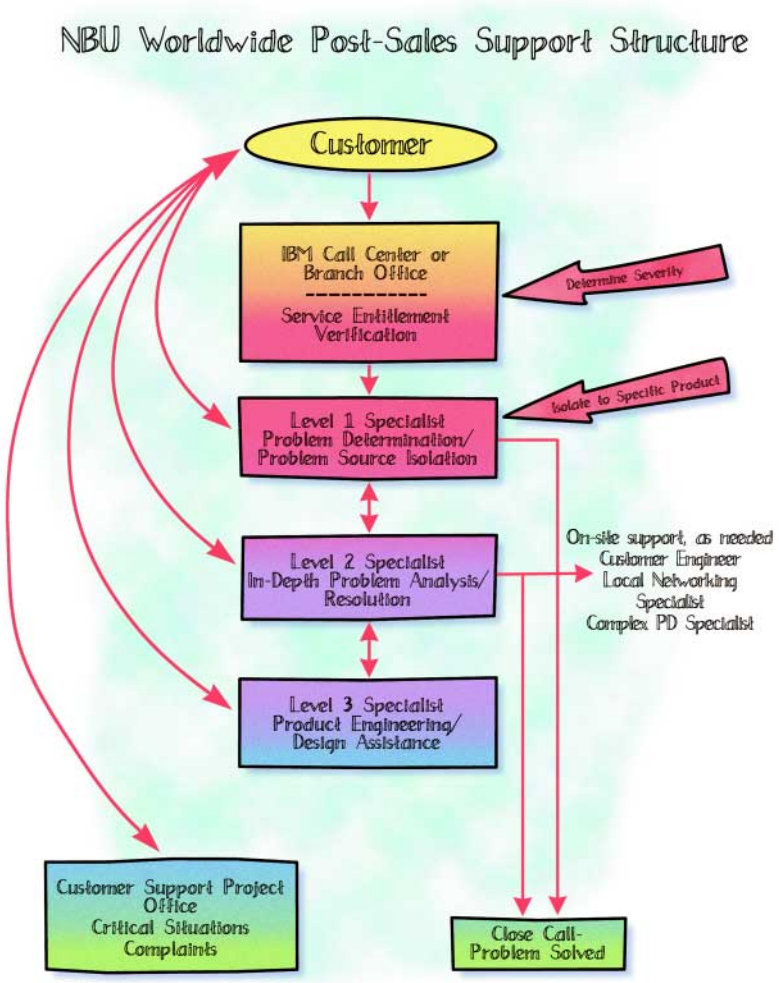

Post-Sales Support

Based on IBM's traditional support model, Product Support Services (PSS), Product Systems Group (PSG), or both, provide Level 1 support when you need post-sales assistance for a Networking Hardware product. The Networking Business Unit provides Level 2 and Level 3/Engineering support when Level 1 support is unable to resolve the problem.

To provide world-class support, we continuously review our Level 2 and Level 3/Engineering support processes, always seeking ways to make improvements. Some of the recent improvements we have implemented include:

• *eSS: electronic Service & Support*

Consistent with IBM's e-business focus, product-specific support information is available on the Internet. For product information, technical tips, FAQs, downloads and forums go to the IBM home page at **ibm.com**, select Support, and then select **Networking** in the Product support homepages pull-down menu, or go directly to: **ibm.com**/networking/support/products.nsf/support/home?OpenDocument

Whether you need the latest code level or technical tip information, or you wish to submit a forum entry, help is accessible directly on the Internet. Information is continually updated in a real-time manner using Lotus® Notes® databases and Domino™ servers. As new information becomes available, support personnel enter it into the applicable database. After being reviewed and approved, the information is automatically made available on the Internet. Our support team is highly motivated to share technical tips and respond to forum entries.

If you haven't already done so, subscribe to Networking Tech Support Web site today. You will automatically receive e-mail notification of new code downloads, tips and FAQs for the products that you specify.

• *PMR Analysis Enhancements*

Product Management Records (PMRs) are used to record and track customer-reported problems. When you report a new problem to IBM, the support specialist (usually Level 1) creates a new PMR in an online database (RETAIN®) that is accessible by IBM support specialists worldwide. This PMR is updated each time an activity is taken to resolve your problem. If it becomes necessary to involve additional IBM skills, other support specialists can quickly review your PMR status, no matter where in the world they may be located. After the problem has been resolved, the PMR is "closed." However, it remains easily accessible in the online database for approximately 30 days before it is archived. During this time, the PMR can be reopened if necessary.

Over the past few months we have significantly enhanced our ability to analyze the current status of PMRs for our products. Every PMR that is referred to an NBU Level 2 support specialist is updated with specific problem information to be captured in a standardized manner. Status codes identify who is responsible for performing the next PMR activity. For example:

a. Level 3/Engineering needs to generate a fix
b. Level 2 needs to analyze a dump or trace, or
c. The customer or service representative needs to provide additional problem information.

Dates are recorded so that we can identify problem areas that may need additional focus. For example, the person responsible for performing the next PMR activity might not have acted in a timely manner.

Now that we can readily identify PMRs that are waiting for NBU action, for example, NBU Level 2 or Level 3/Engineering action, we have stepped up our efforts to ensure that any problems you may experience with our products are addressed in the most expeditious manner.

• *Aged PMR Focus*

Networking products are complex, and sometimes we are not able to resolve your problems as quickly as we'd like. It might be necessary to re-create an illusive problem or capture a network trace. Complex network problems might cause a PMR to remain "open" more than 30 days. We refer to these PMRs as "Aged" PMRs. Over the past 9 months, we have worked diligently to reduce their numbers. By monitoring and taking corrective action where appropriate, we significantly reduced our Aged PMRs during the last half of 1999.

We ran in to a minor setback in 1Q2000 as customers returned to normal network update and expansion activities after their Y2K concerns subsided. However, we are now back on track and we are committed to continuing the significant Aged PMR reduction that we were able to achieve during 1999.

Our goal is to minimize the time required to resolve reported problems.

Rest assured, our name, SSD Networking Business Unit, is new, but our commitment to providing you with solutions, products and services that meet or exceed your expectations remains unchanged.

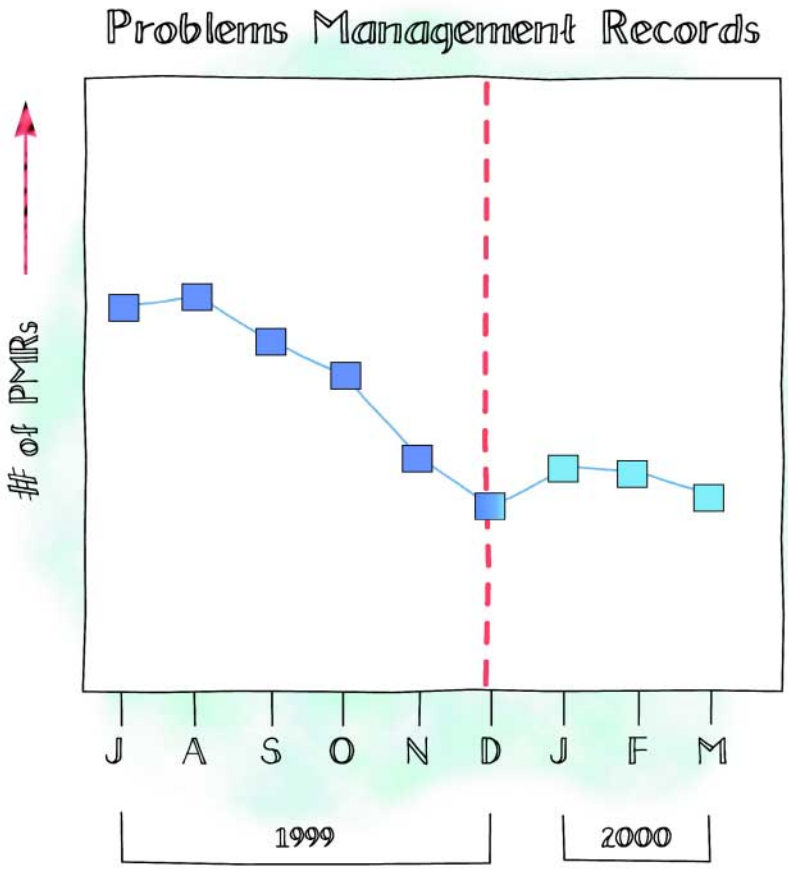

Walt Reston (reston@us.ibm.com) joined IBM as a customer engineer in 1966. Since then he has been a service planner, product planner and since 1995, a product engineer. His current responsibilities include coordination and presentation of the technical support scorecard.

Puzzle

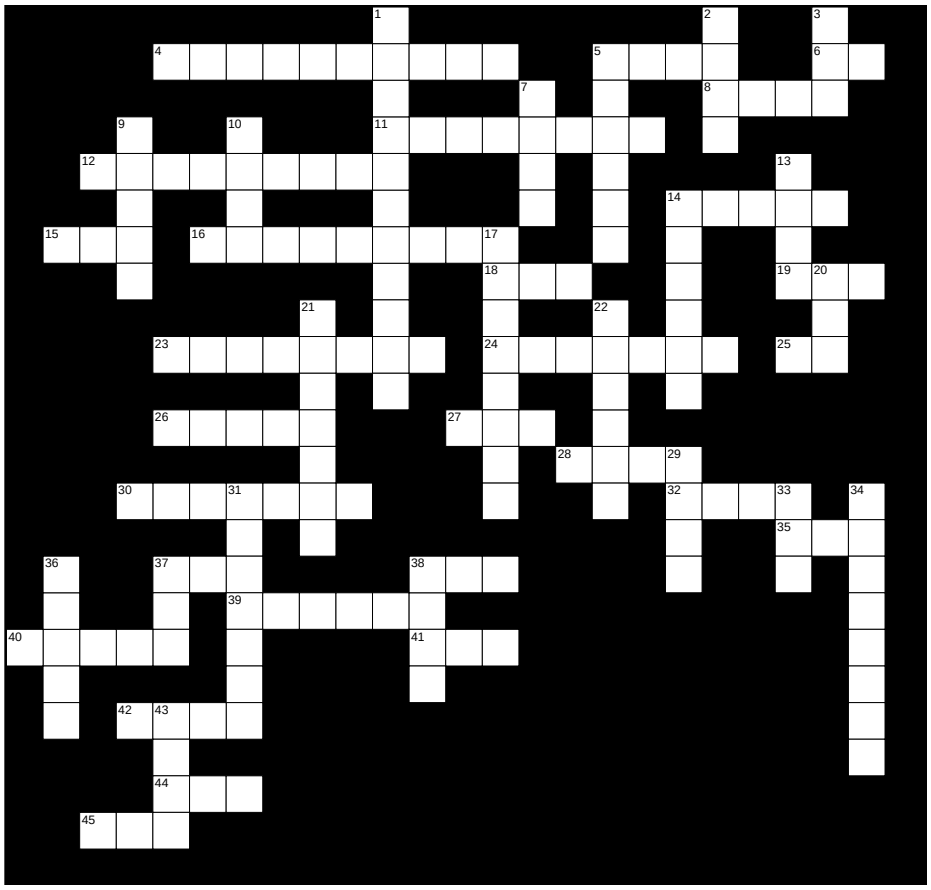

Across

4. Faster in Europe than in the U.S.
5. Billion
6. 28 T1s
8. PU Type 5
11. Monitors NCP
12. Software hidden in hardware
14. Interpreter written by now very rich guy
15. TH, RH, and RU
16. Asynchronous
18. APPN routing scheme
19. Internet names to address systems
23. Named for medium thought to fill space
24. Channel capacity law
25. Delivers TCP
26. Inside the firewall
27. 51.84-Mbps fiber
28. SNA for Peers
30. Must have for NTuneMON
32. Packet Internet Groper
35. 7-layer network model
37. Promised rate
38. e-business company
39. Internet protocol
40. Most-common DCE
41. 56K modem standard
42. Web-page language
44. Southern Silicon Valley
45. Very reliable communications architecture

Down

1. Invented WWW
2. Programming language, coffee
3. Uses 53-byte packets
5. 1 With 100 zeros
7. Million billion
9. Free OS
10. Base of a tree
13. Starts SNA session
14. French engineer for whom signaling speed is named
17. OSI Layer 1
20. PU T4
21. Forerunner of the Internet
22. Famous author, attorney, W.W.I ace, beagle
29. X.25 for NCP
31. Unreal circuit
33. Domain of U.S. government
34. Keeps hackers at bay
36. "Double every 18 months" law
37. Largest domain
38. 28-bit Internet addressing
43. 1000 billion

Down
1. Berners-Lee, 2. Java, 3. ATM, 5. Googol, 7. Peta, 9. Linux, 10. Root, 13. Bind, 14. Baudot, 17. Physical, 20. NCP, 21. ARPANET, 22. Snoopy, 29. NPSI, 31. Virtual, 33. GOV, 34. Firewall, 36. Moore, 37. COM, 38. IPv6, 43. Tera

Across
4. Frame relay, 5. Giga, 6. T3, 8. VTAM, 11. NTuneMON, 12. Microcode, 14. BASIC, 5. PIU, 16. Start-stop, 18. HPR, 19. DNS, 23. Ethernet, 24. Shannon, 25. IP, 26. Intra, 27. OC1, 28. APPN, 30. NetView, 32. PING, 35. OSI, 37. CIR, 38. IBM, 39. TCP/IP, 40. Modem, 41. V90, 42. HTML, 44. RTP, 45. SNA

NCP and 3745/46 Today— Yesterday

Much of the information in *NCP and 3745/46 Today* remains relevant for some time after publication. For that reason, from time to time, our readers request information covered in previous editions. This index addresses that request. The order number of each publication appears beside its period of publication. Each issue also remains available on the Internet on our Web site at:
ibm.com/networking/ncp

Fall '96, G325-3426-04

- Introduction
- IBM Multiprotocol Solutions for the Expanding Enterprise Network
- What's Coming in November 1996? NCP V7R5!
- Network Congestion Control
- STATEMENT and KEYWORD Word Search Puzzle
- IBM Olympic Network
- Library Update
- IBM 2210 Nways: Data Communications Tester's Choice
- ISDN support in NCP V7R5
- Did You Know …
- User group meeting schedule as of July 1996
- IBM 3746: Prime Host Access Device for the Intranet and Internet World
- NetView Performance Monitor - The Good Guy
- What's new in the NTune Family?

Spring '98, G325-3426-05

- About This Issue
- Year 2000 - Ready or Not!
- Network Controller Product Direction
- Extend Your 3746 Investment with the Multiaccess Enclosure
- Library Update
- What's New with ACF/NCP?
- 3745 Token-Ring Connection Balancing
- IP Internal Coupling between NCP and the 3746 Model 900
- Why You Need NTuneMON and NTuneNCP
- What's New in NTuneMON Version 2 Release 4?
- User Group Meeting Schedule as of January 1998
- IBM Outperforms Cisco on the S/390 Channel
- Just How Strategic Is In-Transport Routing?
- 9729 Offers True Optical Networking Technology to Corporate Users
- What Are MTU and MSS?
- First Aid for NCP
- The Redbook Story
- NCP Word Search Puzzle

Fall '98, G325-3426-06

- About This Issue
- 3745/3746 Year 2000 Readiness
- Network Utility Provides IP-SNA Integration and High Session-Processing Capacity
- Networking Implications of S/390 Parallel Sysplex
- Frame Relay/ATM Interworking with IBM's 8265 Nways ATM Switch
- Connectionless Network Prioritization: Is it Real?
- 3174: "Not Just A Display Cluster Controller"
- Focus on 3172 ICP Diagnostics
- Addressing a New Networking Paradigm—High-Speed Connection Networks and Quality of Service
- NCP and 3745/46 Capacity Planning Case Studies
- Network Traffic Analysis (NTA), a Service Offering Since 1989
- What's New with ACF/NCP
- Library Update
- Connection Balancing for 3745 Frame Relay BAN
- 3745 NCP and 3746 Model 900 Supports Switched Frame Relay
- New CIR Support for NCP-Controlled 3746 Model 900 Frame Relay Lines
- NDF Changes Engines
- Emulation Program R14 Now Supports All 3745 Models
- Continued Investment in NTune Family Provides Exciting New

Capabilities

- 3745 and 3746 Model 900 IP Routing
- NCP Service Dates
- Strategic Development of SNA Networks
- Did You Know?
- Do You Use a 3705, 3720, or 3725?
- NCP Word Search Puzzle
- User Group Meeting Schedule as of this Issue

Fall '99, G325-3426-08

- About This Issue
- Meaning of the IBM/Cisco Agreement for SNA Solutions
- Why Is IBM Back in the Cabling Business?
- Reaching Your SNA Host Applications
- Network Traffic Analysis (NTA) Now Accessible through the Web
- Multilayer, Multiprotocol Switched Networking for Powerful e-business and Global Networking
- 3746 Nways Multiprotocol Controller Enhances Major Components in 1999
- TPF adds CDLC/3746 Support for IP
- Library Update
- What's New with ACF/NCP (Carpe Diem)
- User Group Meetings
- NCP Product Set—Year 2000 Analysis Overview
- NCP Service Dates
- ACF/SSP Boasts Enhancements for V4R8
- Connection Balancing for 3746-9x0 Token Ring and Frame-Relay BAN
- Pool Association Change for 3746-900 Duplicate TICs
- Non-Disruptive Route Switching for 3745 Subarea Token-Ring Connections
- NTuneMON Adds Support for Bisync and Start-Stop Lines along with New NCP Functions
- NPSI Update
- NCP Word Search

We want to hear from you...

It has been our goal to make this newsletter informative and interesting.
We hope we have achieved that goal. Please send us your comments and suggestions.

Name _____

Title _____

Address _____

City _____ State _____ Country _____ Zip _____

Fax your comments to:
NCP and 3745/46 Today
c/o Leyland King
NCP Products Development
1 800 426-0121

E-mail your comments via:
NCP and 3745/46 Today at:
ibm.com/networking/ncp
Where you will find an online form on which you can submit your comments or suggestions.

NCP and 3745/46 Today is produced by Department F46A, Hobie V. Love III, NCP Design and Development Mananger

Managing Editor: Leyland King

Technical Editors: G. David Heath
Ralph Nissler

Graphic Designers: Jackie Johnson
Carla Reis

Illustrator: Adam Brill

IBM Corporation
Department TYCA
PO Box 12195
Research Triangle Park, NC 27709
U.S.A.

Printed in the United States of America
5-00

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

The following are trademarks of International Business Machines Corporation in the United States and/or other countries: IBM, the IBM logo, Advanced Peer-to-Peer Networking, AIX, APPN, AS/400, BookManager, DB2, DB2 Universal Database, DFSMSdss, the e-business logo, Enterprise Storage Server, Enterprise Storage Server Specialist, ESCON, FICON, FlashCopy, MVS, Netfinity, NTuneMON, NTuneNCP, Nways, Operating System/2, OS/2, OS/390, OS/400, Parallel Sysplex, RAMAC, RETAIN, RS/6000, S/390, Seascape, SNAP/SHOT, SP2, StorWatch, Sysplex Timer, System/390, VM/ESA, VSE/ESA and VTAM.

Lotus, Lotus Notes and Domino are trademarks of Lotus Development Corporation in the United States or other countries or both.

NetView and Tivoli are trademarks of Tivoli Systems, Inc. in the United States or other countries or both.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft

Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Other company, product, and service names may be trademarks or service marks of others.

Printed on recycled paper

G325-3426-09