NCP/3745 Today

NCP & 3745 Today Newsletter
Networking Controller Products News
Fall 1996 Edition

IBM NCP Products Information Development
IBM Raleigh

Table of Contents

NCP/3745 Today Newsletter

- 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?

Introduction

As you can see in this issue, a lot of new things are coming your way.

The products and features described in the articles provide some of the most important elements necessary to operate a successful intranet, including high reliability, predictability, and performance. We believe that IBM, once again, is proving to be the world leader in the implementation of mission-critical, multiprotocol internetworking solutions.

It is our continued hope that these articles will help you expand your expertise about networking and IBM products.

Again, please remember that your ideas and comments are important to us and we hope you will find this issue as informative and useful as the previous ones.

John Woods,
NCP Product Manager

IBM Multiprotocol Solutions for the Expanding Enterprise Network

Flexible connectivity from desktop to mainframe

June 25 1996 Announcement

Network Computing is all about creating a universally connected world that puts people in touch with each other and businesses, as well as other organizations, in touch with all of their constituencies.

Although network managers are increasingly utilizing the open standards of the Internet and its underlying TCP/IP communications to enable the universally connected world, economic reality requires network owners to leverage existing applications that rely on SNA communications to access critical corporate databases located in S/390 and AS/400 servers.

Thus, today's expanding enterprise network is typically a mix of router-based and controller-based technologies. And, above all else, network managers need network solutions to manage the complexity of today's diverse users, locations, and protocols, while retaining the freedom and flexibility to implement new, cost-effective technologies as they become available.

As the world leader in the implementation of critical, multiprotocol internetworking solutions, IBM is pleased to present the following key enhancements to a selection of enterprise network products:

- Reliable, channel-attached TCP/IP routing on the 3746 platform, providing a complete set of TCP/IP functions, including OSPF-2, BGP-4, and RIP, as well as a rich set of access filters
- High-Performance Routing (HPR) on the 3746 platform to enable you to implement high-availability solutions for both the data center and the network
- Frame Relay support on the 3746 platform, independent of NCP, based on RFC 1490 or boundary access node (BAN) and boundary network node (BNN)
- ISDN support, including new Euro-ISDN primary rate interface (PRI) support of dial access for SNA on the 3746 Model 900 and a new ISDN basic rate interface (BRI) adapter for the 2210 that extends existing ISDN support to additional models

- General functional enhancements of the 2210 to provide enhanced connectivity and availability, additional SNA support, new security features, and more efficient use of network bandwidth
- An additional model of the 2218 Nways Frame Relay Access Device (FRAD) to expand on the recently announced family of low-cost Frame Relay devices

The products and features included in this announcement provide high reliability, predictability, and performance, in addition to appropriate security, backup, and recovery features. These characteristics are some of the most important elements necessary to operate a successful intranet, which is a private network that uses the infrastructure and standards of the Internet and the World Wide Web with firewalls to prevent access from the public Internet. Customers are using intranets to create a business environment in which applications, tools, and information are created, accessed, and shared among different departments, business units, and worldwide offices of a company or organization.

This announcement, coupled with IBM's recent announcement titled "S/390 Network Computing Vision," and other IBM enterprise networking announcements from earlier in the year, consolidates IBM's commitment to help you leverage their existing networking, application, and database investments as you migrate to collaborative computing based on Internet technologies. IBM's Network Computing announcements provide solutions that span the enterprise from the desktop to the largest S/390 server, support secure intranet and Internet access from any desktop, and enable new applications to be easily implemented.

Your needs are driving this new computing model at the same time that technology is making it possible. Corporations want to extend your reach globally, improve quality, cut costs, get to market faster, and serve their customers better. A large, heavy-equipment manufacturing firm that uses SNA and NetBIOS as LAN protocols is one example of a customer who will benefit from this announcement. The hardware configuration consists of IBM mainframes running MVS with CICS and IMS transaction monitors that are connected to PCs running DOS, OS/2, Windows, and Windows NT.

As the customer stated:

> "Our challenge was to provide a flexible alternative that would allow the various business units to access the mainframe data stores and combine that data with whatever types of data and processors they wished to use."

Announcement Highlights

3746 Nways Controller Model 900 and Model 950 Enhancements

Almost all of the following enhancements will be available in September 1996 and December 1996. The enhancements include the following:

- Robust, channel-attached TCP/IP routing that includes a complete set of TCP/IP functions, as well as a rich set of access filters, positioning the 3746 as a key gateway into an intranet.
- HPR connectivity between Parallel Sysplex and the 3746 enables very high SNA application availability, even if there is a host failure. HPR connectivity between the network and the 3746 provides non-disruptive path switching to re-reroute HPR traffic around network outages.

- Frame Relay support that is independent of NCP on the Model 950, and a Euro-ISDN PRI adapter on the Model 900.
- An evolutionary path that allows the 3745/3746 and associated SNA network to evolve and grow in capacity, function, and performance. The 3746 Model 900 is field-upgradeable to a Model 950, and SNA/Subarea applications can evolve to APPN/HPR networking.

See the *IBM 3746: Prime Host Access Device for the Intranet and Internet World* article in this newsletter for more information.

New ACF/NCP V7R5 and *ACF/SSP V4R5*

The new ACF/NCP and ACF/SSP software will be available in November 1996. NCP V7R5 provides the following functions:

- Support for a native ISDN E1 Adapter
- State-of-the-art subarea router solutions
- Excellent APPN capabilities
- High-speed X.25 SNA transmission capabilities
- Comprehensive Frame Relay switching

SSP V4R5 fully supports NCP including ISDN, duplicate TIC, and frame-relay error count support.

See the *What's Coming in November 1996? NCP V7R5!* and *ISDN support in NCP V7R5* articles in this newsletter for more information about the new functions available with these releases.

New NTuneMON Version 2 Release 3

The new NTuneMON software will be available in November 1996. See the *What's new in the NTune Family?* article for a description of the enhancements available with this release.

New ISDN Adapter for 2210 Nways Multiprotocol Router

The ISDN Adapter, available in August 1996, performs the following functions:

- Supports basic rate interface (BRI) ISDN
- Allows greater network design flexibility
- Can be used for either primary network connection or switched backup capability
- Supports a wide variety of signaling specifications

The 2210 BRI ISDN Adapter provides a four-wire, twisted-pair S/T interface with an RJ-45 connector. The ISDN port provides two 64-Kbps B-Channels for data and one 16-Kbps D-Channel for signaling.

The BRI ISDN adapter supports:

- Euro-ISDN
- INS-64

- National ISDN-1 and -2
- AT&T 5ESS
- Nortel DMS-100
- TS 013

New Functions for Nways Multiprotocol Routing Network Services (2210 Software)

The functions for the 2210 Router, available in June 1996 (ISDN August 1996), include the following enhancements:

- WAN re-reroute extensions
- Frame Relay over BRI ISDN dial circuits
- Synchronous Data Link Control (SDLC) secondary support
- Additional IPX capabilities
- X.25 DTE transport
- Point-to-Point Protocol (PPP) data compression
- Data Link Switching (DLSw) support for LAN Network Manager (LNM)
- User Datagram Protocol (UDP) broadcast forwarding
- DECnetV/OSI Echo support (RFC 1139)
- IPX filter enhancements
- Software support for the new 2210 BRI ISDN Adapter
- OSI Echo
- Support for the Australian ISDN switch standard
- PAP/CHAP password security for PPP links
- ISDN Cause Code Management

Nways Multiprotocol Routing Network Services Release 3 Extended includes a range of functional enhancements. Enhancements to the current WAN re-reroute function detect PPP and X.25 link failures (in addition to Frame Relay) and re-reroute traffic over any data link control, including PPP, X.25, and Frame Relay (in addition to ISDN and V.25bis dial circuits).

WAN re-reroute extensions over PPP-, X.25-, and Frame Relay-configured interfaces are enabled through enhanced MRNS WAN support for modems and ISDN terminal adapters that store dial digits.

Additional new functions include Frame Relay over BRI ISDN dial circuits to enable BAN traffic to be transported over ISDN, secondary link stations support for SDLC devices, and software data compression on PPP links.

Additional IPX capabilities include new filtering function and enhancements to the IPX protocol to allow IPX to operate more efficiently over dial-on-demand circuits. Four new IPX filters operate on IPX packets at the interface level rather than at a global level associated with all interfaces.

The X.25 DTE transport function enables the transport of data packets over a TCP/IP network between 2210 routers with attached X.25 DTE devices. Extensions to the DLSw function enable DLSw to transport LNM packets. New UDP support enables the forwarding of UDP broadcasts.

DECnetV/OSI support has been added for Echo support (RFC 1139) to enable Echo requests and replies to be used for the determination of OSI protocol connectivity. The addition of PAP/CHAP password security includes the establishment of passwords at each end of a PPP link and a procedure to check those passwords before bringing up the link.

ISDN cause code management allows user changes to the behavior of the router if a given cause code is generated by the switch.

New 2218 Nways Model 12X Frame Relay Access Device

The 2218 Nways Model 12X Frame Relay Access Device (FRAD), available in June 1996, supports the following functions:

- Introduces low-cost entry-level solution for linking branch offices to a centralized data center.
- Supports two communication ports with no LANs.
- Provides device/host interconnection across a Frame Relay network.
- Channels SNA, BSC, asynchronous, and multiprotocol LAN traffic onto a single Frame Relay permanent virtual circuit, eliminating the need to have separate WAN links.
- Is fully compatible with other IBM products that support RFC 1490.

Fully compliant with RFC 1490, the 2218 Model 12X is designed to ensure interoperability with enterprise network technology today and tomorrow. The 2218 Model 12X:

- Adds another model to the family of 2218 Nways FRADs announced in April 1996.
- Interoperates with IBM 3745 and 3746 controllers, and 3172 software.
- Integrates and supports serial protocols including SDLC, X.25 QLLC, BSC, asynchronous, and SNA downstream physical units.
- Attaches to the host with SLDC, BSC, Token-Ring, or Frame Relay circuit.
- Is fully manageable by NetView and SNMP managers.

Benefits

- HPR provides increased productivity resulting from very short response times and 100% network availability. HPR provides powerful solutions for consolidation of large networks, with reduced operational costs, thanks to maximized line utilization.
- The rich IP support turns the 3746 into a gateway device for a company's intranet.

- Multiprotocol benefits include native routing of IP, APPN/HPR, and SNA protocols over the same media, allowing for link consolidation and reduction of operational costs. Communication with downstream NCPs over Frame Relay BAN connections is enabled, as well as support of IP applications at marginal cost. The PRI ISDN adapter supports the Euro-ISDN standard and operates in conjunction with ACF/NCP, allowing SNA equipment to communicate with ACF/NCP over ISDN connections.

- Management enhancement benefits include support of the IBM 3746 IP router in an SNMP environment, resource monitoring allowing for timely capacity planning, configuration management ensuring the Network Computing availability, and graphical display for IP and APPN/HPR network topologies.

- Advanced Communication Functions and management functions include ACF/NCP Version 7 Release 5, ACF/SSP Version 4 Release 5, and NTuneMON Version 2 Release 3. These software updates provide state-of-the-art subarea router solutions and solid IP support. They are compatible with all ACF/NCP and associated product releases supported by ACF/SSP Version 4 Release 2 still in service. NTuneMON Version 2 Release 3 provides faster resolution of potential and existing internetwork problems, increased network and NCP availability through interactive online tuning, and enhanced management of controllers and networks.

- The announcement of additional 2210 software function and expanded ISDN availability provides enhanced connectivity, improved network availability, expanded SNA support, additional security features, and improved bandwidth utilization. The IBM 2210 family of routers provides an extensive range of connectivity, protocol, and price that can be tailored to link remote offices and regional and central host sites into a cohesive, cost-effective multiprotocol network.

- The 2218 Nways FRAD Model 12X provides multiple protocols over relay transmission links, with priority queuing and bandwidth allocation. Benefits include efficient bandwidth utilization, consistent response times, enhanced session reliability, and simplified network topologies. These allow customers to consolidate networks without mandating change at the network protocol level, and eliminate expensive serial links between the host and the branch office.

- Another benefit to the enterprise network is the ability to combine the Enterprise Network products with the Personal Communications family of products to create total enterprise networking solutions, providing desktop communication across all the major desktop operating systems.

  These solutions provide customers with 3270 and 5250 emulation, SNA client support, and applications for OS/2 and DOS/Windows workstations. They also provide a variety of ways to connect to S/390, AS/400, and S/36 systems, including Token Ring and Ethernet, NetWare for SAA 3270, TCP/IP TN3270 and TN5250, SDLC, dial-in asynchronous communication, Twinax, APPC 5250, console support for AS/400, S/3x, 3270 passthrough via AS/400, coaxial DFT and CUT, APPC 3270, and with the addition of the ATM S/390 Open Systems Adapter 2 (OSA-2).

  The ATM OSA-2 brings to the enterprise LAN environment the security, backup and recovery, availability, and centralized access to data that comes from the S/390 server. When combined with ATM Forum-compliant LAN Emulation (LE) for Ethernet and Token Ring that is also provided by the ATM OSA-2, customers have a powerful addition to their network

that allows them to migrate to ATM without losing their current investment in Token Ring or Ethernet LANs.

Positioning

- These are key products for customers planning to implement IBM's networking strategy for Network Computing (NC) and Switched Virtual Networking (SVN). The IBM 3746, along with the S/390 Open Systems Adapter (OSA), the IBM 3172 Interconnect Controller, and the RS/6000, are IBM solutions to support network peripheral access and routing to S/390 servers in LAN and WAN environments.
- The Nways controllers are the prime server-attached products for high-availability Network Computing in a multiprotocol environment. With HPR support, they enhance Parallel Sysplex environments, allowing fast routing techniques and automatic recovery. Concurrent support of HPR and SNA with APPN/DLUR allows users to evolve at their own pace from SNA to HPR, taking full advantage of and security from the product platform's evolution. Future product evolution includes support of 155-Mbps ATM for LAN and WAN attachment.
- Advanced Communication Functions and management function benefits will suit customers considering migration to a higher level of NCP, customers seeking ISDN PRI E1, customers seeking to expand HPR routes to span multiple APPN subnetworks, and internetwork managers implementing ISDN features of the IBM 3746.
- The 2210 with ISDN can be used at remote locations that require dial-on-demand or dial backup capabilities in addition to the broad array of multiprotocol routing support in the 2210. At a data center location, the 2210 with ISDN provides a very cost-effective, central-site router for ISDN termination.
- The 2218 Model 12X provides a low-cost solution for migrating to networks of the future, including SNA customers who are migrating their branch office networks, customers seeking more cost-effective, faster connections than are currently utilized over multipoint, leased-line networks, and customers seeking to connect smaller sites to corporate Frame Relay networks.

What's Coming in November 1996? NCP V7R5!

By John Foley

On June 25, 1996, IBM announced a new release of the IBM Network Control Program (NCP) in conjunction with the IBM 3746 Models 900 and 950. ACF/NCP Version 7 Release 5 for MVS, VM, and VSE protects your investments by adding new capabilities and finely tuning existing functions in the 3745 and 3746 controller family. Now is the best time to level set your 3745s and 3746 Model 900s with the latest release of NCP.

Our newest release of ACF/NCP Version 7 Release 5 is sure to please you with its enhanced Frame-Relay, Token-Ring, and High Performance Routing (HPR) support along with new ISDN primary rate adapter capabilities conforming to European Euro-ISDN standards.

The main focus of our new release was to implement functions that provide you with strong reasons to migrate to the newest release of NCP while at the same time easing a migration path to the 3746 Model 900.

We concentrated on functions and capabilities that:

- Increase network availability with alternative ISDN backup lines
- Enable duplicate addressing options in token-ring environments for additional backup configurations
- Provide you greater flexibility in your network schemes with additional connection capabilities
- Lower your operational cost by consolidating lines and increasing throughput
- Increase your productivity with dynamic congestion controls
- Protect your investment as we continue to provide enhancements to both Advanced Peer to Peer (APPN) and High Performance Routing (HPR) environments in your SNA backbones thus enabling around the clock availability and nondisruptive service.

These new functions include:

- A native Primary Rate ISDN adapter on the 3746 Model 900 conforming to European EURO-ISDN standards and providing dial-on-demand and backup capabilities for SNA connections flowing over frame-relay links.
- Duplicate token-ring interconnect (TIC) addressing for subarea connections now allow greater addressing flexibility for mixed subarea and peripheral networks.
- Boundary access node (BAN) format for frame-relay subarea links allows remote 3720s, 3725s, COMTEN, and AMDAHL communication controllers connection capabilities to 3745s across frame-relay networks using token-ring attached IBM 2210 routers.
- Dynamic hot standby for multiple link transmission groups (MLTGs) allows you to define a percentage of the total congestion threshold on a multiple link transmission group when a hot standby MLTG is activated to alleviate network congestion.
- Frame-relay tuning enhancements gives a network manager greater abilities to avoid or alleviate network congestion periods with definable parameters to delay information frame transmissions, ignore large bursts of erroneous traffic and ignore subsequent backward explicit congestion notification (BECN) frames for a specified period of time.
- In conjunction with a future release of VTAM, NCP will enable High Performance Routing (HPR) border node capabilities allowing HPR routes to span over multiple APPN networks.
- In addition, NCP enables HPR traffic over X.25 ODLC links on the 3746 Model 900 thus improving performance and availability of these X.25 links.

We are proud of our products and thankful for your continuing support. We hope that this new release of NCP enhances your abilities to provide your customers with the most comprehensive and highest quality networking services available in the world.

NCP Version 7 Release 5 in conjunction with ACF/SSP Version 4 Release 5 will be available on November 29, 1996.

Network Congestion Control

By John G. Waclawsky Ph.D.

With the explosion of networking throughout the world, congestion control is becoming more and more important. Network data rates and volumes are continuously increasing and stressing today's network architectures. Many people are finding that poor congestion control is more catastrophic with higher speed networks, such as ATM, than with lower speed ones.

Congestion control is important for a network to achieve its two primary functions:

- Moving a quantity of data, measured by throughput (or "goodput" might be a better term because retransmissions only reduce throughput).
- Moving data at a desired and consistent speed, measured by response times and the allowed variance in response times.

Congestion control, along with capacity planning and network design, is key to obtaining good throughput from existing network capacity and achieving reasonable response times. In addition, effective congestion control:

- Allows a more stable network environment, because network resources can better tolerate and control traffic spikes.
- Promotes more efficient network operation, because network resources can consistently operate at higher utilizations without the fear of congestion collapse or wasted capacity from retransmissions.

The need for congestion control increases exponentially with increasing network speeds. At gigabit speeds, an enormous number of bits can be lost in a few milliseconds. In poor congestion control environments, link utilization measurements can show high resource use but "goodput" may be very low. Besides the performance problems that arise in a poor congestion control environment, no one wants to pay for retransmissions of the same data multiple times through their network.

This article examines the effectiveness of transmission control protocol (TCP) congestion control and its performance behavior when internet protocol (IP) is the underlying transport. TCP assumes that an unreliable IP transport layer exists. TCP depends on multiple timers to detect errors and to control error recovery. If a packet acknowledgement does not arrive for any reason (for example, a packet or acknowledgement is delayed or discarded), TCP simply retransmits.

What is congestion control?

Congestion control is the ability of a network architecture to effectively control the amount of data allowed to enter a transport. A transport is the collection of the shared nodes, links, and algorithms that are responsible for actually moving data any significant geographic distance.

A transport is a physical entity. The amount of data entering and leaving the transport at any time must be balanced to maintain stability. One of the following rules always governs transport behavior:

- RULE 1: **Data going IN can exceed data coming OUT**
- RULE 2: **Data going IN is always less than or equal to data coming OUT**

There are two primary approaches to congestion control.

- One approach anticipates packet discards and relies exclusively on sender retransmissions when errors and problems occur. This is TCP. With this approach, the transport attempts to minimize queue storage requirements and provides only a ”best effort” to deliver your data, with no guarantees. If storage is limited in the transport, packets are discarded. An example is an IP transport. TCP/IP environments follow RULE 1.
- The other approach, which follows RULE 2, does not rely on sender retransmission. It usually guarantees data delivery by allowing some queueing in the transport and it closely controls traffic flows within and at the entry to the transport. An example of this approach is an SNA or advanced peer-to-peer networking (APPN) network.

In both approaches, window protocols are the basis for congestion control.

What are window protocols?

Today, end-to-end window protocols are the basis for congestion control in both TCP/IP and Subarea SNA networks. Window protocols can allow a receiver to control a sender, which is the underlying basis of flow control. Unfortunately, transport dynamics (routing changes) and packet discards can affect window protocol activity. Window protocols respond to the perceived increase in delay or loss of packets by restricting the flow of traffic into the transport. End-to-end window protocols perform best in deterministic connection network environments. Throughout this article, a simple example demonstrates congestion control concepts and provides insight into TCP/IP congestion control behavior.

An example:

Consider a ”two-hop” route through an existing mesh network between the sender and receiver. This two-hop route will have a window protocol operating between the sender and receiver. Assume this specific route is *one of many* possible paths over which packets flow from the sender to the receiver. Along this particular route, a 512 kilobits per second (kbs) link empties at an intermediate node into a queue feeding a 256 kbs link that ends at the receiver.

Sensing the Bottleneck!

For congestion control to work effectively, it is fundamental that any sender transmitting packets into a transport be able to ”sense the bottleneck”. This is key! Interfering with this fundamental

”sensing” capability severely threatens network stability and degrades performance. From the sender’s perspective, the bottleneck can be found in one of three places:

- A transport element (affected by node processing power or link capacity) is the bottleneck.
- The receiver (affected by processing power or available buffers at the receiver) is the bottleneck.
- The sender itself or the traffic arriving at the sender is the bottleneck.

Senders identify all three kinds of bottlenecks using window protocols.

A transport element is the bottleneck

Node processing power or link capacity could be the bottleneck. A window protocol allows a sender to sense a physical or logical bottleneck in the transport. Assume our sender has a file to transfer. In this case, the sender is not the bottleneck. Consider the two-hop route example.

- First, when a sender is busy transmitting, its window allocation (the sender’s current window size) must deplete. Windows are designed so that the entire window allocation will be exhausted and enter the transport under load. In the example, because physical link capacity is limited and the receiver cannot acknowledge any data it hasn’t seen, the receiver can send acknowledgments to the sender only at a rate of 256 kbs or less. It’s physically impossible for data to get to the receiver or for acknowledgements to be received by the sender any faster! Because the sender can transmit at a rate of 512 kbs and receives acknowledgements at half the 512 kbs rate (256 kbs), the window allocation at the sender (no matter how large) will be depleted over time.
- Second, when the window allocation is exhausted, any additional transmissions from the sender can occur only when an acknowledgement arrives from the receiver. A sender uses the received acknowledgement packets as ”strobe pulses” of network capability. Because the window allocation is depleted, the sender must transmit at a rate that matches the network physical bottleneck capacity (256 kbs in our example) which is a reflection of the rate that data arrives at the receiver. If other flows intersect and share either the 512 kbs or 256 kbs link, data from this particular sender might arrive at the receiver at a rate less than 256 kbs (depending on the severity of cross-traffic interference). In this case, acknowledgements from the receiver will strobe the sender at this newly detected logical bottleneck rate. This logical rate is only the actual capacity of the transport used by the data flow (in our example it will be 256 kbs or less).

As network file transfers, e-mail, voice, video and other traffic types start and stop, logical bottlenecks will form, move, and disappear within the transport. Traffic flows from sender to receiver are likely to converge and compete for limited network resources. When convergence happens, traffic flows interfere with each other (causing queue waiting time), which forces traffic arrivals at each receiver to diminish. The receiver ”slows down” strobe acknowledgements to each sender, because traffic arriving at the receiver has slowed down. The senders, in turn, reduce the amount of data they send into the transport. This behavior is very important. *It keeps the network stable*

by preventing uncontrolled queue growth within the transport. Transport bottleneck sensing is automatic with end-to-end window protocols. This hidden dynamic aspect of congestion control can be found in both TCP/IP and Subarea SNA networks.

Smooth and efficient acknowledgement packet flow is critical to any congestion control scheme based on end-to-end window protocols. The delay (or timing) between acknowledgement packets, which are returning from the receiver to the sender, carry network capability information (or the rate at which packets are processed at the bottleneck). Interfering with acknowledgement packets will severely curtail the effectiveness of end-to-end window congestion control, reduce performance, and destabilize the network environment.

The receiver is the bottleneck

Processing power or the availability of storage at the receiver (outside of the transport) could be the bottleneck. With end-to-end window protocols, the receiver always has the option of strobing the sender at its own (less than transport capacity) processing rate. Packet flow exiting the transport at the receiver matches the receiver's ability to store and process arriving packets. Thus, with window protocols, the receiver can directly control the sender's transmission rate by controlling the acknowledgement rate. This flow control protects the receiver from being overrun by arriving transport packets. *To work properly, the transport must not interfere with acknowledgement packets returning from the receiver to the sender or with the timing between them, because the acknowledgements carry the rate at which the receiver is able to process or store packets*.

The traffic arriving at the sender is the bottleneck

The sender might have insufficient data or processing power to utilize the window allocation. The resulting transmission rate into the transport is less than both the transport and the receiver's capabilities. This condition can exist when an application cannot generate data fast enough or internal contention for physical resources exists at the sender. The end result is that network and receiver capabilities will be underutilized by this particular sender. Many networks experience this state at different times of the day. Hopefully, peak loads or other senders will converge in the transport to use excess capacity. If not, the network was over-engineered with too much capacity and probably costs too much.

Connection, Connectionless and Congestion Control

Connections through a transport are required for good performance and absolutely necessary for good congestion control. *Connections* are logical paths over fixed routes through a transport. Connection networks generally maintain information about the state (or availability) of transport resources. Thus, the transport is controlled, resource usage is controlled, and state information can be used to manage and assign new routes based on resource availability in the transport. Connection networks are usually RULE 2 environments.

In a *connectionless* network there are no logical paths or fixed routes through the transport. Because data flow is not controlled with paths, data packets and acknowledgements can take different routes through the transport. The availability of transport resources can change freely and is

unpredictable. In this stateless environment, the entrance to the transport is not aware of network conditions or the availability of resources for data transmission. RULE 1 governs the transport behavior of these best effort, packet-forwarding networks.

The lack of a path has disadvantages. Connectionless transports can interfere with the operation of end-to-end windowing and therefore congestion control.

The next two sections examine transport performance and congestion control behavior in a connectionless environment. Connectionless environments can cause performance problems and significantly reduce the effectiveness of higher layer congestion control mechanisms.

Performance Problems in a Connectionless Environment

Connection-oriented network environments have significant performance, management, and stability advantages, because a path exists through the transport. Let's now see what kind of performance problems the lack of a path can cause.

Sensing the wrong bottleneck

In a connectionless environment, traffic flow from the sender to the receiver can use a set of transport resources different from returning acknowledgements. Consider the two-hop example. The sender-to-receiver flow traverses a 512 kbs link and then a 256 kbs link. Under ideal conditions, the sender transmission rate should stabilize at 256 kbs (matching the physical bottleneck). If acknowledgements take a different route and return to the sender over a slower resource than 256 kbs (for example, a 56 kbs link resource), acknowledgements can only arrive at the sender at a rate of 56 kbs (this is a physics problem). The sender senses the wrong bottleneck (the 56 kbs link). When this happens, the sender transmits at a 56 kbs rate even though 256 kbs worth of capacity is available. This behavior causes response time problems, throttles traffic throughput, and wastes network capacity. This problem does not occur in connection-oriented networks (such as SNA and APPN) because connections ensure that window protocol traffic uses the same set of resources in either direction.

Lack of preferential treatment for acknowledgements

To obtain good performance between a sender and a receiver, acknowledgements must flow freely through the transport. Even if acknowledgements take the same route (in reverse) through the transport as the sender-to-receiver flow, any queueing of acknowledgements in the transport can seriously impair traffic flow in the sender-to-receiver direction. Queueing of acknowledgements will impact (slow down) the rate that acknowledgements arrive at the sender. The sender will sense increased acknowledgement queue time as the (wrong) bottleneck and slow down. In addition, needless TCP sender retransmission occurs if acknowledgements are delayed too long. Network "goodput" decreases because the TCP sender interprets the lack of an acknowledgement as a lost packet. The end result is unnecessary response time problems, throttled traffic throughput, and wasted network capacity.

Other than the adoption of a connection network, there is no easy solution to this problem. Some vendors offer router prioritization as the answer. But router prioritization is computationally ex-

pensive and *cannot guarantee* that acknowledgements will not wait. Often vendors will suggest that we prioritize on the size of packets but this does not solve the problem. Size prioritization causes all packets of a certain size or less to receive preferential treatment when queued at transport nodes, even those that are not acknowledgements! Also, this approach will not work when acknowledgements are piggybacked on other traffic flows. It is not practical to eliminate the piggybacking of acknowledgements because many small, inefficient packet flows in the transport result, which wastes network capacity. Also, this piggybacking of acknowledgements is the main objective of TCP's "silly window syndrome" avoidance algorithm.

Other options are the assignment of priority levels to specific devices on the basis of their medium access control (MAC) address or the use of some sort of "smart filter". The first option will create very large tables in very large networks. Both options require expensive high skilled personnel to configure and manage the network. Even then, these approaches still cannot guarantee that an acknowledgement won't wait. In addition, router prioritization schemes do not interoperate well and are proprietary solutions.

In a connectionless environment, prioritization is just guesswork and can cause two problems:

- Increased potential for congestion collapse caused by senders retransmitting the *delayed lower priority* traffic.
- Timeouts for lower priority traffic caused by significant delays.

The prioritization problem for connectionless environments continues to be a major issue in computer science research. In addition, if a connectionless environment cannot provide preferential treatment for acknowledgements, how can it provide prioritization for other kinds of traffic (Quality of Service)? These are a few of the reasons that drove ATM to adopt a connection-based transport.

Acknowledgement discards

When a transport node discards a packet, how does it determine that the packet is not an acknowledgement or carrying an acknowledgement? A connectionless environment is stateless and lacks the ability to identify acknowledgement packets (or any sender-to-receiver or receiver-to-sender flows). Any discarded acknowledgements can have a severe performance impact on traffic flowing in the opposite direction in a couple of ways.

- A missing acknowledgement will prevent transmission of additional traffic that is waiting for permission given in the acknowledgement.
- A missing acknowledgement is likely to cause needless retransmission.

In the second case, a TCP sender cannot tell if a packet was discarded. Because the acknowledgements haven't arrived, the only choice for the sender in a connectionless environment is to assume the packet was discarded. (Note: connection networks do not usually discard data. If they do, they usually employ mechanisms that can detect acknowledgements using state information and prevent their discard).

Retransmissions in error recovery

In a connectionless transport, such as IP, error correction is not considered a requirement. It is simply assumed to be the responsibility of a higher layer. If any packet is corrupted, it is simply discarded. This philosophy can have a significant (as well as confusing) performance impact in a noisy environment when a transport node elects to have a packet wait while the node attempts retransmission or recovers a temporarily out-of-service resource.

In a connection network, traffic flow to the problem resource simply slows down and ultimately stops without retransmissions. Consider what happens in a connectionless network when traffic is delayed for any reason (such as, node or link error recovery activity). Acknowledgement timers pop and senders retransmit packets. The load at the troubled resource increases and could precipitate a congestion collapse. *A good congestion control scheme should not introduce new traffic to the transport during periods of congestion*. Rerouting performed by the transport can theoretically help in this situation. In practice, however, routing changes are performed too infrequently to be of much help and the problem becomes severe with increased network speeds.

Sender retransmissions are the wrong solution when congestion control (slowing or stopping all packet transmission) is required. Retransmission at the wrong time can become the equivalent of throwing gasoline on a fire to put it out! The fundamental problem is that the sender in a connectionless environment cannot tell the difference between a *lost* and a *delayed* packet. In fact, all the connectionless transport performance problems discussed above become more severe with increased network speeds. *These points support the conclusion that connectionless environments are not suitable as high speed transports*.

Congestion Control Limitations in a Connectionless Environment

Connectionless networks also have severe congestion control limitations due to the very nature of the connectionless philosophy. *Three questions must be answered accurately for congestion control to work effectively in any network*:

- How is congestion detected?
- Which packet do you throw away (or delay in a queue)?
- Where is congestion reported?

The dynamics and stateless nature of connectionless networks make effective congestion control extremely difficult (if not impossible) to achieve. In a connection environment, a network path, which facilitates the detection and reporting of congestion, makes it possible for:

- A connection transport to report about the availability of transport resources to the traffic origination points at the network boundary.
- A sender to slow down or stop transmission if problems occur and not have to worry about discards and retransmissions.

Compare what happens in a connection environment versus a connectionless environment when we answer each of the questions.

Detecting congestion

Usually high buffer usage indicates congestion. In connectionless environments, users typically share a common storage area. This sharing provides the opportunity for one misbehaving user to affect the transport performance seen by others. Any mischief can also be compounded by retransmission activity. In a connection transport, the ability to identify connections (and if necessary constructing a queue per connection) prevents one abusive user from affecting the service level of others. It helps detect congestion and, if necessary, pinpoint the individual, misbehaving user.

Which packet do you throw away?

The primary question here is: What action does a node in the transport take when it becomes congested? With a stateless environment, the origin of traffic arriving at a congested network node is not known. Nodes could examine the packet header, but this extra processing cannot be performed in a high-speed environment. Typically, in a connectionless environment, discarding packets is the usual response to congestion. The decision of which packet(s) to discard is important. When responding to congestion, should a node discard an arriving packet or an existing packet in the queue? If a node discards a packet from the queue, how does it decide which one? IP networks do not have state information to determine the cause of congestion and the best recovery actions.

The fundamental philosophy of the connectionless environments is that there is no state; therefore, congestion control is problematic. We certainly do not want to throw away high priority traffic. Prioritization is most important when congestion is occurring. Without an effective transport-wide congestion control scheme, the network is unable to fairly prioritize different types of traffic. In a connectionless transport, detecting and dealing with congestion is extremely difficult.

Reporting congestion

How is congestion reflected to the senders that are causing the problem? TCP/IP does have an optional feature called source quench. A transport node can issue a source quench packet to indicate congestion. However, because an IP node doesn't know who is causing the congestion, *there is no guarantee that the location that receives the source quench packet caused the congestion problem.* If a TCP sender is **not causing** the congestion problem and elects to slow down because of source quench notification, it only penalizes itself and makes even more bandwidth available to the traffic causing the congestion. Also, while the source quench packet is making its way through the network, the transport could reroute traffic making the source quench packet irrelevant. For these reasons, most IP network products do not support source quench activity. Any rerouting decision is also uncontrolled (that is, not planned or managed to ensure overall network stability or to meet Quality of Service objectives). Load may be simply transferred elsewhere, potentially causing thrashing or other problems.

A partial solution?

One potential solution to some of the congestion control limitations unfortunately does not solve all of our connectionless performance problems. The solution is to prevent limitations from becoming

an issue by using a design and capacity planning strategy that deliberately underutilizes transport components. That's why router advocates respond to performance problems with three words: "add more capacity". In fact, this line of reasoning implies that overcapacity is the only effective congestion control mechanism for connectionless environments. This philosophy not only wastes money but has limited effectiveness in bursty, high speed environments. Also, it is well known that carriers make more money and users save more money when network resources are effectively used.

Think Strategically: There is a fundamental difference!

Many popular protocols (such as, TCP and Subarea SNA) have performance limitations due to their use of end-to-end windowing algorithms. In addition, the very connectionless nature of IP is a serious problem for effective window and congestion control activity by TCP (or any other protocol carried by IP for that matter). Please be aware that APPN offers a connection-oriented network solution that does not use end-to-end window protocols and eases many of the window limitations that exist today. Window protocol limitations and performance problems are eliminated entirely by IBM's High Performance Routing (HPR) algorithms which provide a windowless rate-based congestion and flow control environment.

Conclusion: Cheap vs Inexpensive networks

The attractiveness of end-to-end window-based congestion control techniques decreases with increasing transmission speeds and geographical size. Further investigation (for example, into the costliness of the transport storage needed to support end-to-end window protocols) will reveal that end-to-end window mechanisms are legacy technology that is nearing the end of its useful life for connectionless as well as connection networks. These mechanisms have served us well in the past but, regrettably, must be replaced with high speed networks to function effectively and efficiently. Connectionless protocols are simply old technology from another networking era, dominated by low speed WAN connections.

The connectionless philosophy of IP causes too many window protocol limitations and congestion control problems, which exponentially grow with increasing network speed, to be considered seriously as a high speed transport. Window protocol restrictions can be significantly reduced by using a "hop-by-hop" implementation (for example, APPN) and totally eliminated in the HPR windowless rate-based design. (Note: a hop-by-hop window approach is not possible in an IP environment because a connection through the transport is required for hop-by-hop window schemes to work.) APPN today can support link speeds well into the hundreds of megabits per second range. HPR is the ultimate solution because it provides a windowless rate-based flow control environment (like ATM) and can scale to the gigabit networks of the future.

With HPR, SNA has migrated beyond the old limited window-based network environments to a platform that is capable of productively using high speed network components. Because IP interferes with window protocol operations, perhaps the answer is to unhook TCP from IP, spoof the window acknowledgements, and send traffic over a manageable connection transport, such as HPR and ATM. HPR and ATM can use a high speed transport, provide QOS (Quality Of Service), and adequately support asynchronous traffic flows. Make wise technology investments. Don't be fooled! What looks like an "inexpensive network today" based on old connectionless protocols

such as IP might be a ”cheap” network tomorrow.... and tomorrow will be here very soon!

STATEMENT and KEYWORD Word Search Puzzle

Directions

Find the statements and keywords used in NCP generation for frame relay in the puzzle below. The words are formed in the diagram forward, backward, up, down, and diagonally, but they are always in a straight line and they are never formed by skipping over any letters. Letters may be used more than once in the diagram, and the names often overlap.

Note: All the statements and keywords can be found in the NCP, SSP, EP Resource Definition Reference (SC31-6224-04) that goes with NCP V7R5, SSP V4R4, and ER R12.

```
S R E M I T E C S R S Y H P E
E T A R O L U D O M A X G Z M
R E M R B L L O L C K G O I A
R O P C K C O L B L N U T E R
O B I N I R O D B I P I O R F
R E L A Y C N A K R T T P P X
T N S M A I T C D S L S X A A
O A I P W A O E R L E N A R M
T T N N D L S T U O X A M X S
L A Y D C R E M O T T O L M I
A D A T F M A X P U C O M D A
C C O S S E R D D A D D A T R
O P A T A D X A M R E S P O L
L
```

Word List

```
ADDR          FRSEDRPU      NPACOLL
ADDRESS       LMI           PHYSRSC
ANS           LOCADD        PORTADD
BLOCK         LOCALTO       REMOTTO
CLOCKING      LSPRI         SPOLL
COMRATE       MAXDATA       SRT
DATABLK       MAXFRAME      SUBPORTS
DLCI          MAXOUT        TIMERS
DYNWIND       MAXPU         TITIMER
ERRORT        MODULO        XMONLNK
FRELAY
```

IBM Olympic Network

By Jim Ashton

IBM, the worldwide information technology sponsor of the 1996 Summer Olympics, provided a client/server implementation with networks that delivered information to all facets of Olympic production and to spectators in Atlanta and around the world. This super system of networks spanned over eighty locations, supported hundreds of applications, and encompassed all IBM platforms.

Despite initial setbacks, a results system tracked and reported the results of all 271 medal events to officials, spectators, and the world media and to the Olympic World Wide Web site. Information systems provided official results, athlete biographies, and history for commentators as well as event schedules, and local weather and restaurant information for spectators. Other support systems included attendee ticketing, parking and security, housing, transportation, personnel, and administration.

The master Results Systems database resided on a S/390 server and contained information about the athletes and officials, competition start lists, world and Olympic records and results from in-progress and completed competitions. Information was downloaded as needed to each venue, where it was stored in the PS/2 Results Server DB2/2 database in the sports venue. As the competition took place, scoring, timing, distance, and other statistical information was entered through IBM PS/2 Results Workstations, located in the venue, and then sent to the PS/2 Results Server. The Results information was managed by IBM's DB2 database software on both the workstation and the central server.

The real strength of the client/server implementation was a system that spread the tasks and enhanced operations at each geographic location for maximum performance and availability at the field of play and the venue operations site. At the same time everything was networked so that information was shared and available to all points that needed it.

Network Components

With high availability and reliability, the networks spanned more than 80 venues (from Stone Mountain to the Olympic Stadium to the Georgia Dome to Wolf Creek to Savannah and on and on) with 25 front-end processors (FEPs), 25 2217 Nways concentrators, 450 bridges, and 600 hubs, 2,000 miles of fiber, and 5,450 miles of copper. They interconnected four mainframes, 80 midrange systems and 7,000 PC workstations. The Results System network itself connected PS/2s, S/390 mainframes, AS/400s, and RS/6000s with 250 token-ring networks. Interoperability enabled the Results System to seamlessly interface with Swatch timing devices to calculate event results and disseminate the information to the world.

The overall information-delivery structure consisted of a technology operation center, data center, backup data center, and the triangle consolidation area. Their locations were determined by the availability of buildings and the needs of the components. The primary data center (mainframes, AS/400s, and lots of DASD) occupied a sprawling low-rise building in the suburbs. The data center had two 3746 61A Model 900s. Located in downtown Atlanta, the technical support center, with its thousands of people, connected to the data center over miles of cable. The backup data center stayed mobile in trailers downtown.

The greatest number of people and equipment, known as the triangle consolidation area, congre-

gated primarily in three locations downtown in what was referred to as the ”Triangle”, or affectionately the ”Bermuda Triangle”. The INFORUM and Atlanta Merchandise Mart buildings housed the press and media center and the technology operations center. The other two triangle corners were the Georgia Technology campus including the Olympic village and the Georgia World Congress center, which was a large sports venue. Each of the three sites used 3745s, 2220s, and many 8260s (hubs). Each had a 3746 Model 31A 900 and NCP Multinetwork Server (NMS) and there were 400 hubs scattered around the three locations. The 2220s in each location were connected to each other through a private T3 frame-relay network.

The networks for sports venues were categorized by size based on the number of events, people, and necessary equipment. For example, Moorehouse, one of the basketball sites, was considered a small venue because it only housed one event. Stone Mountain stadium, a medium venue, housed the marathon, cycling, and tennis.

There were a total of twenty 3745 Model 170s (most with an attached NMS) throughout the medium and large venues. A medium sports venue network consisted of a 3745, an NMS, a 2217, an AS/400 and 8229s on an apex token ring connecting applications for results, commentator information, distribution, and spectator information systems, each running on separate token-ring networks. A large sports venue had these components plus a second AS/400 and applications for its own operations system. The 3745 connected to the data center and to two triangle corners through leased public frame-relay circuits. Also an 8229 connected to the data center through ISDN lines.

The small sports venue network consisted of a 2217, an AS/400, and 8229s connecting applications for results, distribution, and spectator information systems, each running on a token-ring network. The 2217 connected to the data center through leased public frame-relay lines and an 8229 connected through ISDN lines.

For application data, DB/2 was used on all IBM platforms (OS/2 DB/2, RIS6000 DB/2, and AS/400 DB2).

Network Technologies

The 1996 networks took advantage of advanced peer-to-peer networking (APPN), token-ring, and frame-relay technologies. Broadcasting networks are primarily analog and the only digitalized photos were for photo finishes provided by Swiss Timing. These still photos each required about 1.3 megabytes of data and could be sent in less than a second on a 16-megabit token-ring network. Therefore, the bandwidth needed did not require ATM technology.

APPN and frame-relay technologies enabled the networks to deliver with critical high performance. The frame-relay implementation consisted of combined public and private frame-relay networks with:

- Instantaneous switching to backup links
- Seamless end-to-end management
- Dynamic bandwidth management

APPN enabled integration of SNA networks and interconnection between multiple platforms.

The NCP Multinetwork Server (NMS) enabled both subarea and APPN capabilities for each NCP. The 3172 Model 390 NMS provided distributed network management by allowing small portions of the network to run autonomously enabling local error recovery and security. For example, the accreditation application network needed to be autonomous because badge access required constant availability.

Commercial Relevance

Ninety percent of the applications from sponsors and all the hardware used in the Olympic network were off-the-shelf and are currently available. The Olympics production, however temporary, is similar to a big billion dollar commercial business. It is a practical model to replace small, medium, and large venue networks' applications with business-oriented ones for insurance, banking, transportation, and other business needs all with 3745s connected to an S/390.

Library Update

By Bob Yelvington

We told everyone who received the Spring 1996 edition of this newsletter about the removal of the version and release numbers from the covers of the manuals that were revised with NCP Version 7 Release 4. In an effort to reach all our readers, we are repeating the information below.

When you migrate to a newer release of NCP, you may notice that some of the documentation that is shipped with new releases has not been revised. Most of the manuals in the library are revised for each release. Those manuals that are not revised are still shipped with the new release to ensure that you have the entire documentation package. The version and release information will be in the Edition Notice in the front matter of each manual.

There have been small changes in the information in two manuals in the library. These manuals are not being revised for this release of NCP. They are the *NCP X.25 Planning and Installation* manual and the *NCP, SSP, and EP Generation and Loading Guide*. The changes have been included in Chapter two of the *NCP Version 7 Release 5 Library Directory* under the chapter section where each manual is discussed.

Supported Releases of NCP, SSP, and EP

The supported releases of the NCP, SSP, and EP are listed below to help keep you informed.

Product	Release	Operating Systems	End of Service
NCP	V4R3.1	MVS, VM, VSE	Not Announced
	V5R4	MVS, VM, VSE	Not Announced
	V6R3	MVS	Not Announced
	V7R1	MVS, VM, VSE	12/97 (MVS), 10/98 (VM)
	V7R2	MVS	10/98

Continued on next page

Product	Release	Operating Systems	End of Service
	V7R3	MVS, VM	Not Announced
	V7R4	MVS, VM, VSE	Not Announced
	V7R5	MVS, VM, VSE	Not Announced
SSP	V3R6	VSE	Not Announced
	V3R8	VM	
	V3R9	MVS	Not Announced
	V4R1	VM, VSE	10/98 (VM) VSE Not Announced
	V4R2	MVS	10/98
	V4R3	MVS, VM	Not Announced
	V4R4	MVS, VM, VSE	Not Announced
	V4R5	MVS, VM, VSE	Not Announced
EP	R3	VSE	
	R4	MVS, VM	
	R6.1	MVS, VM, VSE	
	R7	VSE	
	R8	MVS, VM, VSE	
	R9	MVS, VM, VSE	
	R10	MVS, VM	
	R11	MVS, VM	
	R12	MVS, VM, VSE	

World Wide Web

NCP information is available on the World Wide Web. From the IBM Networking home page (`http://www.raleigh.ibm.com`), you can link to a list of products and choose Software and then choose the Network Control Program home page. You can go directly to the NCP home page by entering `http://www/raleigh.ibm.com/375/375prod.html`. (This is the fastest approach.) From the NCP home page, you can view:

- NCP and hardware announcements
- NCP Library Directory
- NCP Migration Guide
- and this newsletter.

IBM 2210 Nways: Data Communications Tester's Choice

The IBM 2210 Nways Multiprotocol Router earned a Data Communications Magazine Tester's Choice in their July testing of routers with data link switching (DLSw). According to the *DLSw Routers: Scaling to New Heights* article in the July 1996 issue of Data Communications:

> The IBM 2210 Nways Multiprotocol Router fielded nearly 3,000 SNA sessions in our tests (and may be capable of handling yet higher levels in live production networks), while turning in stellar results in our prioritization tests. The big plus is price. At $5,545, it's the least expensive router we tested.
>
> All we can say is that the 3,000-session figure shouldn't be taken as a absolute maximum. The 2210 also turned in nearly flawless scores in our prioritization tests, coming closer than any other router to hitting our target bandwidth utilization levels.

For more information about the testing, read the published results posted to the Data Communications web site at `http://www.data.com/Lab_Tests`.

The 2210 has also demonstrated success where it counts the most – in customers' networks. Worldwide customer acceptance of the 2210 in the first half of 1996 has been almost as high as the total for 1995: EMEA led in terms of total customer acceptances, while US attainment exceeded US customer acceptances for all of 1995.

The IBM 2210 family of routers provides an extensive range of connectivity, protocols and prices that can be tailored to link a customer's remote locations and central host sites into a cohesive, cost-effective multiprotocol network. The 2210 provides connectivity for users of TCP/IP, Novell's NetWare (IPX), SNA, DECnet, AppleTalk, Banyan VINES, and Net BIOS.

The IBM 2210 Nways Multiprotocol Router was enhanced by the recent announcement of an ISDN adapter for four high-end 2210 models and a broad range of new software capabilities.

For the most current information about the 2210, surf to:

`http://www.raleigh.ibm.com/220/220prod.html`

ISDN support in NCP V7R5

By Rick Konitshek

NCP V7R5, in conjunction with the 3746 Model 900, now provides a native primary rate integrated services digital network (ISDN) adapter. This adapter conforms to the European Euro-ISDN standards.

A new line interface coupler, known as Line Interface Coupler-16 (LIC16), is controlled by a Communications Line Processor (CLP) that can also support synchronous data link control (SDLC), Frame Relay, and X.25. This primary rate interface includes support for one D-channel operating at 64 Kbps and 30 individual B-channels each operating at 64 Kbps. The D-channel is reserved for call signaling and the B-channels are reserved for all data flows.

NCP implemented ISDN support for SNA connectivity using frame-relay framing over ISDN. It allows connectivity with remote SNA stations that also support frame-relay framing over ISDN or are attached to a router, such as the IBM 2210. With this support, ISDN provides yet another connectivity option for SNA sessions.

NCP V7R5 ISDN has been implemented as a dial-on-demand (DOD) function. When NCP receives a session activation request, it will automatically 'dial up' a new B-channel connection to the network if there was not one already connected that could be shared. When there are no

longer any sessions using a B-channel connection, the B-channel connection will be dropped after a user-specifiable amount of time.

Backup Capability

The DOD support provides the flexibility to use ISDN for backup purposes. Because NCP's ISDN function includes full SNA support, the backup capability is not limited to frame-relay links. ISDN can be used to provide backup for any media type, attached to either the 3745 or 3746 Model 900, that supports SNA sessions today.

However, ISDN is primarily intended for backup of SNA sessions on frame-relay links. If such a session fails, a C-list or an operator command can establish the failed session over an ISDN connection. This backup can be done by connecting directly to the remote SNA station or router over the ISDN.

DOD capability also allows you to add more bandwidth or backup capability to a multilink transmission group (MLTG) when needed. ISDN subarea links can be included in NCP's mixed media multilink TGs (MMMLTG) and can be specified as hot-standby links.

System Definition

NCP represents the ISDN interface with a non-switched D-channel line and PU1 definition. Each of the 30 B-channels are represented with a switched B-channel line and PU1 definition. These B-channel definitions are associated with a particular interface (D-channel definition). Because NCP performs frame-relay framing over ISDN, the B-channel line is analogous to a frame-relay physical line and the B-channel PU1 is analogous to the frame-relay local management interface (LMI) PU. NCP represents the sessions with logical lines and PU2, PU2.1, or PU4 stations. These logical line and PU definitions are pooled and can be activated over any ISDN interface. Each B-channel connection can be thought of as an individual physical 64 Kbps frame-relay line. Each has its own LMI support and supports the full range of DLCIs. Most definition parameters that apply to Frame Relay, now apply to ISDN as well.

HPR

High Performance Routing (HPR) is fully supported for ISDN connections. Error recovery procedures (ERP) or non-ERP mode can be used. Automatic Network Routing (ANR) from one ISDN link to another ISDN link is outboard. The term 'outboard' means that the 3746 Model 900 routes the frame from one link to another without the involvement of NCP. This technique improves performance. ANR routing between an ISDN link and a 3746 Model 900 frame-relay link is also outboard.

NPM

NCP V7R5 ISDN will include Network Performance Monitor (NPM) support. Performance data can be collected for B-channel connections (B-channel lines) as well as at the individual physical

unit level (logical lines). The performance data counters will be the same as those for frame-relay physical and logical lines. Also, with an ISDN accounting function, NCP sends NPM records indicating when a session starts and ends. This information will be useful for calculating B-channel connection durations. The existing boundary session accounting function can provide information about the ISDN SNA logical unit sessions.

NtuneMON

NTuneMON supports ISDN resources. NTuneMON V2R3 includes several new panels to display information about ISDN resources. The panels display all the relevant SNA and frame-relay information similar to frame-relay resources. In addition, the panels display important new ISDN information, for example, the party numbers and subaddresses. See the *What's new in the NTune Family?* article in this newsletter for more information.

ACF/TAP

ACF/TAP provides a new function specifically related to ISDN. ACF/TAP and NCP enable you to perform a trace of the entire ISDN interface. This trace captures the data flowing on all B-channels as well as the signaling information flowing on the D-channel. You can also perform a trace on only the D-channel, a single B-channel, or on a single logical line. A new ACF/TAP option enables you to segregate line trace data for a specific B-channel from a trace of the entire ISDN interface. This option will come in handy when the early stages of the B-channel connection initialization need to be traced. Because you might not know which NCP B-channel definition will be used for the connection, you can trace the entire interface and then pull the information for the specific B-channel from the formatted trace.

Did You Know ...

By Ann Bunn and Joe Fowler

A service update package (SUP) for NCP V7R4 became generally available on July 19, 1996. This SUP includes maintenance at a service level equivalent with PUT tape 9604.

A SUP for SSP V4R4 became generally available on July 19, 1996. This SUP includes maintenance at a service level equivalent with PUT tape 9605.

The following APARs should be installed to use X.25 on the 3746 Model 900 (NCP V7R4):

- IR32053
- IR32055
- IR32056
- IR32059
- IR32061

- IR32375

Refer to DOC APAR IR32039 for updates to *X.25 Planning and Installation* (SC30-3470-10).

The following HIPER APARs should be installed:

- IR31997
- IR32305

All the APARs are available on the NCP V7R4 SUP, except IR32375, which is available through ISMD.

User group meeting schedule as of July 1996

Networks Expo	Dallas	Oct. 29-31
Comdex	Las Vegas	November 18-22 1996
Guide	San Antonio	November 17-21 1996
COMNET	Washington, DC	January 1997
BPEC	San Diego	Beginning of February 1997
Networks Expo	Boston	Middle of February 1997
SHARE	Anaheim	March 1997
NETWORLD INTEROP	Las Vegas	April 1997
IBM Tech Conf	Nashville	April 1997

IBM 3746: Prime Host Access Device for the Intranet and Internet World

By Thierry Fredon

On June 25, 1996, IBM announced a new release of the IBM 3746 Models 900 and 950 that dramatically enriches its networking capabilities in terms of multiprotocol support, capacity, availability, and network management. These additional enhancements continue to make the 3746 prime network access equipment in a multiprotocol intranet and Internet world.

Multiprotocol Capabilities

The multiprotocol capabilities of the new IBM 3746 Models 900 and 950 provide network dynamics while keeping all the benefits of SNA. This diversity is delivered with the new support for High Performance Routing (HPR) including Automatic Network Routing (ANR), Rapid Transport Protocol (RTP) (previewed for December 1996) for SNA/APPN traffic, and with a native Internet Protocol (IP) Routing, while providing the full benefits from frame-relay lines and DLCIs sharings and Euro-ISDN backup connectivity.

HPR with ANR operates over ESCON channels, token-ring and Ethernet LANs, frame-relay and SDLC links and provides significant performance improvements. Data throughput for token-ring to ESCON traffic on a 3746 Model 950 HPR configuration can be up to 30 times better than on a 3745 configuration. The IBM 3746 Models 950 and 900 Network Node natively routes IP traffic over ESCON adapter, token-ring and Ethernet LANs, PPP (Point to Point Protocol) leased lines, and frame-relay links.

Frame Relay provides the single multiprotocol transport network solution to the 3746. It complies to RFC 1490 and allows a very large variety of sharing at the line or DLCI level. For example, it supports the Frame Relay Boundary Access Node function, which enables multiplexed communication between remote routers and their downstream physical units. Also, a primary ISDN adapter supporting the Euro-ISDN standard and operating with NCP will allow SNA equipment to communicate with NCP over 30 B channels (30 * 64 Kbps) ISDN connections simultaneously. frame-relay multiprotocol framing will be used to transport the data over the ISDN channels.

Increased Availability and Connectivity

The ability to use dual network node processors (NNPs) enables increased availability. You can configure a second Network Node Processor (NNP) in the 3746 Models 950 and 900 Network Node for back-up. If the active NNP is interrupted, the stand-by NNP can take over, either automatically or under operator control while the established sessions remain active. You can also install a second expansion enclosure to increase the maximum number of processors from 10 to 16 in the 3746 base frame.

Network Management

Network management is extended to include the newly supported protocols and functions with NetView and NPM support. NetView provides you with powerful network management capabilities, including:

- End-to-end management
- Centralized reporting of all network problems
- Network topology display.

NetView for MVS displays information for resources attached to an APPN/HPR backbone network. These resources include SNA dependent physical units (PU type 1 and 2 such as a 3270) through the dependent logical unit requestor and server functions, APPN resources, and HPR (PU type 2.1) resources. Netview for AIX displays information for resources attached to an IP backbone network. NPM reports activity data for the resources used in APPN and HPR traffic while the Controller Configuration and Management application (CCM) handle the box configuration.

The 3746 and Parallel Sysplex

An HPR environment provides MNPS (Multi Node Persistant Sessions) for users end to end. Each VTAM system in the parallel Sysplex is capable of providing the HPR end point function. This is

known as the Rapid Transport Protocol (RTP) function. APPN End Nodes and Network Nodes can be the RTP end points - this is the case for the 3746 Model 950 and 900 Network Node. With the HPR function integrated into the IBM 3746 Nways Controller, you are extending the HPR benefits to your PRIME HOST ACCESS device and so to your network.

NetView Performance Monitor - The Good Guy

By Scot MacLellan

There's one out there now, lurking in some dark corner of your network. It's hiding, waiting patiently for the right moment to strike. It's waiting to quickly take you by surprise or slowly sneak up on you. Disruptive and annoying, and at times downright damaging, this is a frustratingly elusive creature.

What is it?

It's a problem - a performance problem - just waiting to happen.

Performance problems come in many shapes and sizes. There are big, cumbersome ones that creep up slowly, hauling down busy lines and nodes with their sheer weight. Others are fast and sleek, rushing in unexpectedly to cripple small components of your network. All are unpleasant. All are damaging. And all are avoidable.

Avoidable? But how?

Well, wherever there are bad guys, there are usually good guys too, right? The good guy in the battle against performance problems is the NetView Performance Monitor, good old NPM. He's not new - he's been repelling the advances of performance problems for years now - and he's constantly rearming himself to keep up with the latest technology, and to eliminate even the newest performance problems.

There are different ways to tackle performance problems. Sudden ambushes can be countered by a lightning-speed reaction and the immediate pin-pointing of enemy positions. NPM is very good at these tactics. To stop slow, heavy performance problems before they cause damage, NPM works tirelessly in the background, gathering data and building up a dossier that you can use at the earliest stages.

Pin-Pointing

PerfDesk is one of NPM's latest and most effective weapons against these performance enemies, allowing a clear and simple display of the exact data necessary to identify at a glance the source of a performance problem. Using the PerfDesk facility, you can create customized windows containing configuration and performance data for your network, and you can save these views for reuse at any time in the future. When you save windows in a PerfDesk, all the actions required to produce those windows are saved too. PerfDesks can be created by the most senior systems

programmers, whose skill is then reuseable at the click of a mouse button by even the most inexperienced network operator.

Let's imagine that you have a group of users in London that connect to a crucial business application on your host. These users might be connected with TIC3 to a 3746 Model 900 at a remote CCU, which itself is connected to a local CCU by a fast subarea line. One of these users experiences poor response time. How can you quickly pin-point the source of the problem?

Well, from the list of PerfDesks saved on your OS/2 workstation, you click on the one named 'London'. Suddenly, six simple graphs pop up on your screen:

- A pie chart, containing the host component and the network component of the overall response time.
- A bar chart showing the status of each of the virtual routes that can be used for the users' sessions.
- A line graph detailing the CCU utilization of the local 3745.
- A line graph showing the subarea line utilization
- A bar graph showing the utilization of the TIC3
- A line graph showing the utilization of the LAN segment where the users' workstations are attached.

NPM automatically starts collecting this data when you click on the PerfDesk and updates these graphs every 25 seconds. Your first glance, however, tells you that the network portion of response time is unusually high (and a slow application on the host isn't causing the problem). There are no blocked virtual routes, the 3745 is not excessively loaded, and the subarea line is running at only 50 percent of its capacity. You see that TIC3 utilization looks normal, but LAN segment utilization is running at 100 percent! You don't need to waste any more time looking for the location of the problem; it's in the LAN.

Problem pin-pointed in a few seconds. Not bad!

Trend-spotting

The PerfDesk application pin-pointed a sudden ambush in the remote LAN. However, not all performance problems are so sudden. Some sneak in gradually, ever so slowly, hoping that you won't notice them until it's too late. The 'London' PerfDesk showed the subarea line utilization at only 50 percent. Nothing to worry about there, right?

Well, maybe not. But wasn't that same line running at 40 percent capacity just two weeks ago? And 30 percent just last month? Who knows? Who remembers that sort of thing? Who cares? As long as everything is working fine today, there's no need to worry just yet. You can deal with performance problems when they arrive, right?

Wrong. Once some performance problems get their foot in the door, it's already too late.

If you wait until performance is suffering from an overloaded subarea line, there might not be a quick solution. How long does it take to get a new leased line installed or to upgrade your current

line to a faster one? You need to prevent this problem - you can't leave your users suffering for weeks on end while you wait for the installation of a new leased line!

This is a good time to take advantage of the NPM Batch Reporting Program. To get a summary of your network usage, you only need to run a few jobs against the data that NPM has been faithfully collecting each week. It's an early warning system. If you see traffic on your subarea line increasing five percent every week, you can react before it becomes a problem. You can either discover the cause of all the extra traffic (and take steps to reduce it) or two months from now you might need to order an upgrade to that line.

With NPM's batch reports, you can see trends in your network and react before they become problems. You can also study the data to understand where tuning could improve performance and efficiency. Sample batch reports are provided for all of NPM's data collection functions from VTAM buffer pool data to NetWare server performance. You can easily customize existing reports or create new reports to clearly show the data that is most important to you, in your network.

Look mum, no hands

You can use NPM's data to improve your network performance, by tuning or careful balancing. But all this takes time, and perhaps you're too busy with the day-to-day management of your network to think about potential problems. Do you have to keep an eye on the graphs to make sure that everything is running well within capacity and that there aren't any potential performance problems on the horizon? You don't want to wait for an angry phone call before knowing that something's wrong, do you?

Nah. Let NPM take care of all that stuff for you. Don't waste time monitoring the current situation when you could spend that time more constructively. NPM informs you in plenty of time if there is something that requires your attention.

How?

Simple. Set the upper and lower thresholds for each data type that you consider crucial. NPM sends a message or alerts NetView if the values reach these limits. NPM requests your attention only when it's necessary. "Help, there are a load of retransmissions on the frame-relay line to Paris - come and have a look before things get serious!"

This capability can be used creatively as a labor-saving and money-saving device. Suppose your leased line to San Francisco only comes close to full capacity for a few hours a day, during large file transfers. The rest of the time only a trickle of interactive traffic occurs. Why not replace the big expensive fast line with a cheaper line that can comfortably handle the light interactive traffic? When the file transfers start, and the line utilization climbs above, say, eighty percent, an alert to NetView can trigger an automated entry that activates a switched line in the same transmission group. This line can be deactivated later, by a similar method, when utilization falls below a certain threshold. You pay for the switched line only for a few hours each day.

A Guard-dog for our Networking Neighbors

Alert and attentive, NPM keeps a constant eye on your network, ready to bark if any of its neighbors are threatened. NCP and 3746 Model 900 come under its protection, supplying performance data for careful scrutiny. VTAM too, is now guarded by NPM, which gathers performance data to make

sure everything is running fine and keeps an eye on all the sessions - even the most remote - connected to VTAM applications. LAN bridges, LAN segments, and NetWare resources are also under the supervision of NPM.

What's new in the NTune Family?

By Arvis Bridges

No more fumbling around in the dark to find problems. No more trial and error resource allocation. NTune highlights important fields, calls attention to problem areas, and provides online technical information and guidance. NTune helps you monitor your NCP networks, and to interactively change many of your NCP generation parameters.

NTune consists of two products- NTuneMON and NTuneNCP. Together, these two products help you monitor and tune your active NCPs without interrupting your operation.

A new release of NTuneMON and its feature NTuneNCP supports the functions provided with NCP V7R5 enabling you to monitor ISDN resources, modify new congestion control and error threshold suboperands and addressing keywords, and view slowdown data.

NTune now displays information about ISDN physical lines, B-channels, and logical stations. You can see a list of the physical lines defined for a selected NCP including the relative line address, LINE definition label, and status of each line. From this ISDN Physical Line List panel, you can request more detailed information about a selected physical line, such as line speed, physical station details, and line timers. On another new panel, you can see a list of B-channels for a selected ISDN physical line including the address, LINE definition label, and status of each B-channel. More detailed information about a selected B-channel includes party numbers and addresses and B-channel station details. NTune also provides a list of logical stations attached to the B-channels and details about a selected logical station.

To help with congestion control, you can now view and modify two new DYNWIND keyword suboperands for 3745 frame-relay physical lines. DYNWIND(dwd) indicates the amount of time NCP should wait before transmitting the next I-frame for a frame-relay station experiencing congestion. DYNWIND(twc) indicates whether NCP should ignore BECNs for 100 milliseconds after an initial BECN is received.

NTune also displays new tuning fields for certain types of connections. You can view the source service access point (SSAP) for logical token-ring connections on the ATUTL Token-Ring SNA Station panel for all logical stations regardless of whether they are subarea or peripheral. You can change the addressing for token-ring subarea stations and the locally administered MAC address for the physical line. You can modify both the source and destination service access point for a logical line.

The ATUFP Frame-Relay Physical Line Details panel now displays the error count threshold for 3745 frame-relay physical lines and the current number of errors. You can modify the value of the error count threshold at any time regardless of the state of the physical line or station.

NTune now displays the first failure data capture (FFDC) information when NCP enters slowdown. On the new Buffer Slowdown Information panels (ATUBS and ATUBX), you can view buffer status, performance collection status and statistics, and total boundary sessions as well as inbound and

outbound data information for stations and sessions when NCP enters slowdown.

International Business Machines Corporation
Department CBJA
Post Office Box 12195
Research Triangle Park, NC 27709
USA

APPN, AS/400, AIX, CICS, DB2, DB2/2, ESCON, IBM, IMS, MVS, MVS/ESA, MVS/XA, Netview, NTune, NTuneMON, NTuneNCP, Nways, OS/2, Parallel Sysplex, PerfDesk, PS/2, RS/6000, SAA, S/390, VM/ESA, and VTAM are trademarks of International Business Machines Corporation.

Windows is a trademark or registered trademark of Microsoft Corporation. All other company, product, and service names may be trademarks or service marks of others.

© Copyright International Business Machines Corporation, 1996.

Printed on recycled paper.
Printed in the U.S.A.