3TECH

The 3Com Technical Journal

3Com®

Volume 5, Number 1 *January 1994*

High-Performance Scalable Networking with Routed ATM
By John Hart

Enhancing the Performance of Today's LAN Technologies

Transcend Network Management Solution
By Lionel de Maine and Bob Weder

Integrated Management Based on Connectivity Systems and SmartAgent Software

3Com Leads the Multivendor APPN Pack
By Howard Blevins

Implementing IBM's Advanced Peer-to-Peer Networking Code on the NETBuilder II Platform

3TECH The 3Com Technical Journal

3TECH, 3Com's technical journal, is published quarterly by 3Com Corporation, Santa Clara, California 95052.

Officers: Eric A. Benhamou, president and chief executive officer; Bob Finocchio, executive vice president, network systems operations; Christopher B. Paisley, vice president and chief financial officer; Debra Engel, vice president, corporate services; John Hart, vice president and chief technical officer; Doug Spreng, vice president and general manager, network adapter division; Alan Kessler, vice president, systems sales–North America; Janice Roberts, vice president, central marketing; Ralph Godfrey, vice president, volume sales–Americas; Richard Joyce, vice president, sales–Europe and ICON.

Circulation and Policy

3TECH is sent free to qualified subscribers. Send qualification card (enclosed in this issue) and address changes to: *3TECH*, P.O. Box 869038, Plano, Texas 75086-9870 or fax them to 408-764-6477. For customer service questions, call 408-764-6626.

Submissions

Manuscript submissions, inquiries, and all other correspondence should be addressed to *3TECH*'s editor: Suzanne Dowling, 3Com Corporation, P.O. Box 58145, Santa Clara, California 95052-8145. Articles in *3TECH* are primarily authored by 3Com employees; however, non-3Com authors are encouraged to submit for publication articles dealing with 3Com-related research or solutions to technical problems.

ISSN 10519637

3TECH The 3Com Technical Journal

Volume 5, Number 1 • January 1994

Feature Articles

Departments

Inside:

Features

Departments

Announcements

Editor's Note

Focus on HPSN, Transcend, and APPN

By Suzanne Dowling

Suzanne Dowling is managing editor of 3TECH. *Suzanne has also managed the newsletters* SNA Perspective *and* Network Management Perspective. *Before coming to 3Com in 1989, Suzanne worked at Communications Solutions, Inc. and SRI International in writing and editing positions.*

She received her B.A. from the University of California, Los Angeles and has M.A. degrees from the University of California, Berkeley and the University of California, Los Angeles.

This issue of *3TECH* focuses on 3Com's high-speed networking strategy—what it is, how it works, and how managers can use it to improve network performance; on Transcend, the company's innovative network management solution; and on the significant accomplishments 3Com has achieved in the APPN arena.

High-Performance Scalable Networking with Routed ATM

High-performance scalable networking is 3Com's strategy for providing high-speed links to handle the demands of growing networks and advanced network applications. The strategy directly addresses the primary challenge facing network planners today—the ever-increasing need for bandwidth. In this article (pages 3–10), John Hart describes what 3Com's strategy is, how it works, and how managers can use it to improve network performance as they need it without having to resort to large-scale, high-risk "forklift" equipment upgrades.

Transcend Network Management Solution

Transcend, 3Com's new family of network management software, represents a significant advance in simplified management of LAN and WAN networks across major network management platforms. Lionel de Maine and Bob Weder describe (pages 12–19) Transcend's connectivity systems and SmartAgent concepts, and show how these technologies can solve the challenge of managing complex networks.

3Com Leads the Multivendor APPN Pack

The multivendor APPN demonstration at Interop in March 1993 was the latest milestone in a nine-month development effort by 3Com's IBM internetworking team. The APPN development effort continues as IBM releases more robust versions of their APPN code. In this article Howard Blevins describes (pages 23–28) the engineering issues involved in porting the first implementation of the IBM APPN code to 3Com's NETBuilder II platform.

New Bug Reports Department

This issue introduces Bug Reports, a new department that lists 3Com's product software and hardware bugs. The listing is arranged by product family and is organized in an easy-to-read, problem-and-solution format that is consistent across all product lines.

Two New Faxback Services

Starting with this issue of *3TECH*, you will be able to grade each issue you receive by filling out the *3TECH* scorecard on the inside back cover. Simply fax your scorecard back to us—let us know what you think!

And now when you have a change of address or want to be removed from our mailing list, you can fax the subscriber information form to us (see back cover). Thanks!

Does Your *3TECH* Disappear?

If your *3TECH* disappears, we've now got what you need—a subscription card to pass on to your coworkers so they can start getting their own copies (and leave yours alone!). Just have them fill in the card and mail it in today.

Feature Article

High-Performance Scalable Networking with Routed ATM

Enhancing the Performance of Today's LAN Technologies

By John Hart

The continuing demand for more bandwidth in data networks calls for an approach that is both technically sound and responsive to long-term business needs. 3Com's High-Performance Scalable Networking (HPSN) strategy directly addresses requirements with a balanced approach.

HPSN builds on customers' existing hub and router configurations to deliver the speed of asynchronous transfer mode (ATM) and the control of routing, while leveraging existing LAN technologies. It is a balanced strategy that places functionality in routers, hubs, and adapters to best meet customers' needs.

This article focuses on the migration to high-speed data communications within the building and describes what 3Com's strategy is, how it works, and how managers can use it to improve network performance as they need it without high-risk "forklift" equipment upgrades.

Special thanks to Dr. Bob Klessig for his contributions to this article.

Planning for Bandwidth Demand

A network manager planning to meet bandwidth demand by expanding an existing network or designing a new one faces some critical choices. Which high-speed technology will best meet network requirements now and in the future? How should this technology be implemented—with mixed-media hubs, switches, multiplexors, routers, or a combination of products? Can higher bandwidth be provided for backbones, servers, and desktop devices without introducing unnecessary complexity and expense? And most importantly, can this change occur without disrupting the existing network? And how will users of new technology communicate with existing LAN users?

3Com followed specific customer guidelines to meet requirements expressed by network managers from diverse organizations. HPSN is designed to:

- Deploy only in the parts of the network where it is needed
- Evolve as a series of steps to provide significant benefits at minimal incremental cost
- Have each new capability build on earlier enhancements without loss of functionality
- Maintain connectivity throughout the configuration

High-Performance Scalable Networking Strategy

3Com's HPSN strategy provides a step-by-step migration to a high-speed environment. This environment can accommodate various high-speed technologies including fiber distributed data interface (FDDI), 100-Mbps Ethernet, and asynchronous transfer mode (ATM).

3Com's HPSN approach applies to all portions of the network: workgroup, building infrastructure, campus backbone, and wide-area backbone. The focus of this article is on the building backbone as a means of enhancing the performance of the existing LANs on the floor. The migration process involves innovations added to 3Com's existing hub and bridge/router platforms, including the NETBuilder II®, the LinkBuilder® 3GH (third-generation) hub, LinkBuilder MSH™ multi-services hub, and the LinkBuilder FMS™ stackable hubs.

John Hart *is vice president and chief technical officer at 3Com Corporation. He is responsible for providing strategic engineering direction and development for new products and technologies. In addition, he is part of 3Com's executive committee, which manages 3Com's daily operations.*

Before joining 3Com, John was vice president of engineering and advanced development at Vitalink Communications Corporation. While at Vitalink, he invented remote bridging, was awarded several patents, and managed the development of bridges, routers, and network management products.

Prior to Vitalink, John worked for Control Data Corporation, South Central Bell, and Southern Bell in senior network architecture, research, and product development positions.

Throughout his career, John has participated in various network standards activities such as the IEEE 802 committee, and has been a voting member since 1987.

Glossary

ATM
Asynchronous transfer mode; a high-speed cell-switching and multiplexing technology

ATM backbone
A 155-Mbps vertical segment containing one or more ATM virtual channels

Bridge-per-port
A capability that provides full LAN bandwidth to each hub port and its attached end system or systems

Cell
An ATM packet with a fixed length of 53 bytes (48 bytes for payload and 5 bytes for the header)

CellBuilder
Modules that implement an ATM downlink by converting conventional LAN traffic to ATM traffic, and vice versa

LAN segmentation
Dividing LAN bandwidth into multiple independent LANs to prevent performance bottlenecks

Today's Collapsed Backbone Architecture

When going from a single LAN per building to a single LAN per floor, most building networks today are designed with a collapsed backbone routing approach. A collapsed backbone router configuration avoids having a router on each floor by repeating each floor's horizontal LAN segments across a vertical downlink to a single router (typically in the basement and located with a server farm). The collapsed backbone uses a fiber infrastructure from the initial single repeated backbone LAN. This is typically a star configuration with UTP on each floor to a hub, and with the hubs interconnecting vertically through one of the fiber pairs in a fiber bundle containing 12 FOIRL fiber pairs. The server farm LAN segments are directly attached to the single router. This collapses the backbone onto the high-speed backplane of a router. As a result, data in the NETBuilder II bridge/router moves approximately 80 times faster than a distributed Ethernet backbone or eight times faster than FDDI. Hubs in each floor's wiring closet continue to concentrate the LAN segments, but networking intelligence and complexity now reside in the basement with the collapsed backbone router.

Figure 1. *Collapsed Backbone with Multiple LAN Segments on the Third Floor*

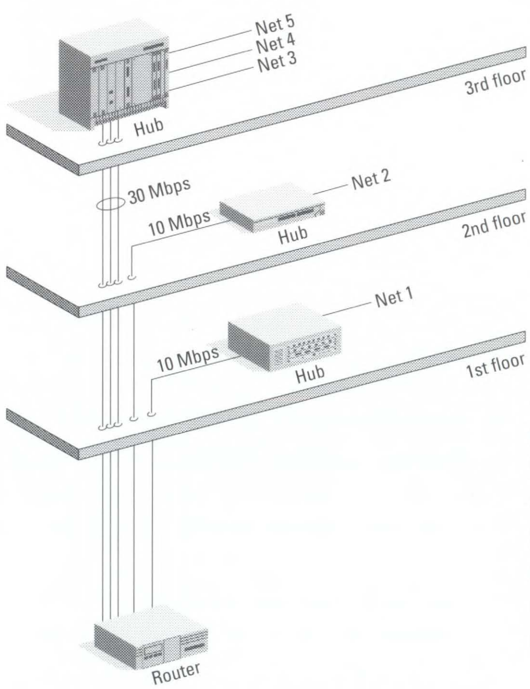

Three horizontal LAN segments are deployed on the third floor, increasing floor bandwidth by a factor of three.

3Com's HPSN strategy uses the collapsed backbone model for a three-step integration of higher performance at a minimum cost.

HPSN Stage 1: Enhancing the Collapsed Backbone with More Horizontal and Vertical LANs

As bandwidth demand grows, a performance bottleneck can quickly result if all the users on one floor share a single LAN. A collapsed backbone can split users across multiple LANs because each new horizontal LAN segment extends vertically to a collapsed backbone router port using a separate fiber downlink. This effectively scales the bandwidth of the vertical infrastructure in proportion to the number of horizontal LAN segments. Figure 1 shows a configuration with three horizontal LAN segments deployed on the third floor.

Additional vertical downlinks don't usually require new cabling installation, since most buildings have spare fiber pairs in each floor's bundle. One constraint on the addition of downlinks is the router's port capacity. The manager must make sure that the collapsed backbone device has enough ports to handle both current and near-term future downlinks from the floors. Of course, multiple collapsed backbone routers can always be located together.

Another constraining factor is increased complexity. Segmentation improves performance, but it also results in additional LANs. For example, if IP is used as a network layer protocol, each new segment requires its own IP network number, complicating management and depleting the organization's number allotment.

To reduce complexity, the manager could group the downlinks associated with the three third-floor segments and connect them to a bridge, which is in turn connected to a port on the router. Bridging the three grouped LAN segments in Figure 1 into one logical workgroup requires only a single IP network number, and the router firewalls this group of LANs from the others.

However, this solution requires an extra bridge, which adds to the delay in the vertical infrastructure and, unless the port connecting the external bridge to the router is equivalent in speed to the three downlinks, congestion in the bridge is also possible. Also, adding a high-speed link between the external bridge and router adds cost.

Port Grouping Creates Virtual LANs /Workgroups

These issues are resolved when a port grouping feature is added to the collapsed backbone router that provides the bridging functionality internally between the three downlinks. Since the "port" to the router is now internal, there is no additional delay, and its speed is proportional to the speed of the three downlinks. The three grouped LAN segments are termed a virtual LAN. The virtual workgroup is defined by the collection of end systems attached to the grouped LAN segments. A virtual LAN that takes up more than one port on the bridge/router looks like a single LAN to the network. But because the administrator can still route traffic between virtual LANs, port grouping retains the advantages of full multiprotocol routing. Further, since port grouping is provided in a router, techniques like proxy address resolution protocols (ARPs) can reduce broadcast/multicast traffic within the virtual LAN.

Figure 2 shows an example of multiple virtual workgroups. In this example, the engineering, marketing, and finance groups are kept separate to isolate data resources and manage traffic between these virtual LANs.

Figure 2. *Example of Workgroups Within a Building*

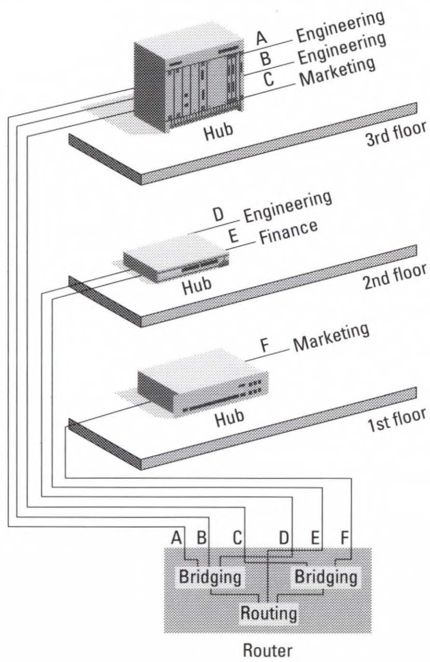

Using port grouping, all the engineering LAN segments are combined into a single virtual workgroup, even though they are physically divided into three segments spread across two floors. The virtual workgroup is assigned a single IP network number instead of three.

Creating virtual workgroups using a single IP network number (because they are on the same virtual LAN) mitigates the complexity of segmentation on the floors (no end system address changes), and makes it easier to manage.

HPSN Stage 2: High-Speed Downlinks Increase Bandwidth and Reduce Segmentation

Increasing bandwidth within the workgroups means increasing LAN segmentation on the floors. But eventually the administrator runs out of spare fiber cabling in the building riser or uses up all the physical ports on the router. This dilemma is resolved by using a single high-speed downlink to replace multiple slower LAN segment downlinks. Overall network performance can continue to be

Glossary (continued)

Port density
The number of ports, physical or logical, per network device

Port grouping
The ability to group ports on a router into subsets; bridging among the ports in the subset, and connecting between the subsets using routing

Route caching
Storage of forwarding information (based on network topology and routing policy) associated with a destination, starting when the first packet to the destination is processed, to speed the forwarding of all subsequent packets to the same destination

Routed ATM
3Com innovations that combine the forwarding speed and low latency of ATM with the traffic control of routing

Virtual workgroup
Any set of physical or logical ports, grouped and treated in bridging and routing as ports on a single LAN

Abbreviations and Acronyms

ARP
Address resolution protocol

ATM
Asynchronous transfer mode

FDDI
Fiber distributed data interface

FOIRL
Fiber optic interrepeater link

LLC
Logical link control

SNAP
Subnetwork access protocol

TCP/IP
Transmission control protocol/internet protocol

increased through further horizontal segmentation without the need to change any hardware or software at the desktop.

Of course, per-port frame processing performance of the collapsed backbone router must increase by a factor of ten to support high-speed multiple downlinks. There are numerous solutions to this performance scaling problem; all require distributing some level of routing functionality within the collapsed backbone router.

Route Caching Provides Scalable Routing Capacity

The preferred solution distributes only the simpler high-performance frame forwarding logic to the port interface cards (termed port switching engines), and centralizes the complex route determination logic in a central routing engine. This approach is termed "advise-and-consent" because the first time a destination is seen by a port switching engine, the central routing engine determines the route and tells the switching engine how to forward subsequent frames with the same destination. The port switching engine then performs the forwarding operation with the advice-and-consent of the central routing engine. The port switching engine remembers the routing information in a "route cache."

Route caching uses many of the caching principles established in mainframe virtual memory caching schemes, but with one significant advantage: each switching engine is responsible for routing only the frames from end systems associated with the attached downlinks, so each port switching engine sees only a few routes relative to the total routes seen by the central routing engine. Also, from the perspective of the port switching engine, the routes the end systems use change infrequently. Both attributes mean that a switching engine's cache hit ratio over a 24-hour period is likely to be very close to 100 percent. With route caching, the frame-forwarding capacity scales proportionally to the number of high-speed downlinks.

Local Boundary Routing Provides Scalable Network Capacity

High-speed downlink support also requires some level of internetworking on the floor to attach Ethernet and Token Ring LAN segments. The challenge is to move simple low-cost internetworking functionality called local Boundary Routing™ (a form of LAN switching) to the floor without giving up the full-function benefits of a collapsed backbone router. The requirements for full internetworking functionality, simplicity, and low cost are the same as those of 3Com's approach to connecting remote offices.

3Com addressed these requirements with its Boundary Routing system architecture by extending the functionality of the collapsed backbone router to the remote site through the extension of the advise-and-consent technique. Only simple frame-forwarding logic is distributed to the remote site. Full internetworking functionality is supplied to the remote site by the central routing engine in the linked router in the central site.

Local Boundary Routing is another extension of the advise-and-consent technique. It behaves like a NETBuilder II port switching engine for its attached LAN segments. Like the port switching engine, if it knows the route, it handles the frame forwarding independently with the advice-and consent of the central routing engine in the collapsed backbone router.

High-Speed Downlinking with LinkBuilder 3GH and NETBuilder II

An example of a unique high-speed Stage 2 downlink solution is shown in Figure 3. Using the LinkBuilder 3GH, Ethernet LAN segments are interfaced to one of three FDDI segments through bridge-per-port technology. High-end servers are attached to the FDDI segments associated with their primary workgroup using either bridge-per-port technology or FDDI concentration.

The NETBuilder II provides full-function routing between the three FDDI segments and a FDDI campus backbone. The result is three extremely high-speed workgroups within a building with complete firewalling between them and the campus backbone. The configuration's performance is scaled by distributing the high-performance intraworkgroup frame forwarding to the LinkBuilder 3GH, while keeping the much more complex route determination logic centralized in the NETBuilder II.

Today, FDDI is the only standard high-speed LAN technology that can be deployed as a high-speed downlink, interconnecting either Ethernet or Token Ring LAN segments. In the near future, 100-Mbps Ethernet could also be used to interconnect Ethernet LAN segments.

An issue with LAN downlinks is that all the segments switched onto the downlink must use the same network number or must have a full-function router at both ends. Of course, multiple downlinks can mitigate this, but it could eventually present a problem as the number of LAN segments increases.

Figure 3. *Three High-Speed Workgroups*

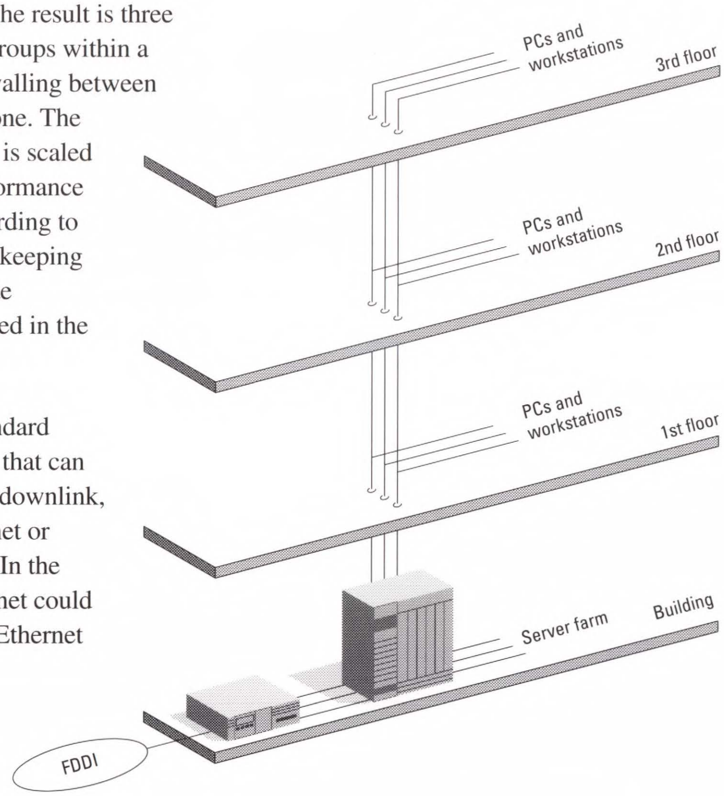

Ethernet LAN segments are interfaced to one of three FDDI segments through the LinkBuilder 3GH. The NETBuilder II provides full-function routing between the three FDDI segments and an FDDI campus backbone. In this example, the NETBuilder II scales the configuration's performance by distributing the high-performance intraworkgroup frame forwarding to the LinkBuilder 3GH.

High-Speed Downlinking with ATM

The 155-Mbps multimode fiber interface specified by the ATM Forum is an ideal downlink technology. ATM offers a number of advantages to managers looking for a high-bandwidth downlink technology to handle advanced network applications and future growth.

Because the identity of the individual LAN segments can be retained by mapping them to individual virtual channels, a single ATM downlink can forward frames from LAN segments associated with multiple network numbers. ATM allows for considerably more LAN segmentation without using up fiber cabling pairs. A single ATM link can easily support 15 to 30 Ethernet or 10 to 20 16-Mbps Token Ring LANs on each floor. Since each LAN segment is mapped to a different virtual channel within the downlink, every segment can be identified by the router. Since the NETBuilder II router can now perform virtual channel grouping instead of port grouping, an administrator can create virtual LANs just as if each segment had its own downlink. The ATM downlink, which uses existing multimode fiber-optic cabling in the building riser, can be implemented just by adding new modules in the collapsed backbone router and floor hubs.

Figure 4. *CellBuilder Support of LAN Segments on an ATM Downlink*

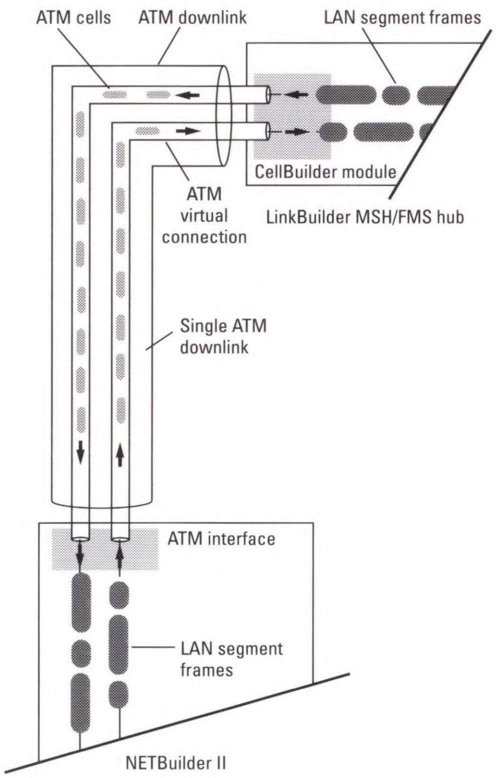

At each end of the ATM downlink, LAN segments are identified from the virtual channel identifier in the ATM cell header. This ensures proper handling of frames and allows low latency cut-through forwarding techniques.

LAN downlinks transmit variable-sized frames and consequently have variable delay, known as latency. ATM downlinks overcome variable latency by segmenting frames into short, fixed-length blocks called cells. Data, real-time voice, and video transmissions can all be combined with this approach.

Once frames have been segmented into cells, the router can also significantly reduce its latency. Since all of the routing information is normally contained in the first cell, the frame-forwarding decision does not have to wait until all the cells have been received, especially if the destination port is also a ATM interface. 3Com calls this cut-through routing technique stream routing, and when it is used with route caching, it significantly reduces the latency of the data network.

CellBuilder Technology for ATM Downlinks

3Com supports ATM downlinks using CellBuilder™ technology within LinkBuilder MSH and FMS hubs to convert Ethernet and Token Ring frames sent across the ATM downlink into ATM cells. When cells are received from the ATM downlink, the CellBuilder technology performs the reverse process, reassembling the ATM cells into LAN frames for transmission to the stations. Figure 4 illustrates the conversion between LAN frames and ATM cells on the building backbone.

CellBuilder uses the CCITT standard ATM adaptation AAL5 protocol for ATM segmentation and reassembly. To identify the type of LAN, the module uses the logical link control (LLC) protocol and the subnetwork access protocol (SNAP) standardized by the Internet Engineering Task Force in RFC 1483.

HPSN Stage 3: Collapsed Backbone with Routed ATM

The standards-compliant ATM downlink just described can be connected to an ATM switch (Figure 5), which gives an administrator the ability to direct (cross-connect) each virtual channel, and therefore each LAN segment in a specific router or router port.

Cross-connecting LAN segments improves performance because the traffic load is shared across routers as well as across "homing" LAN segments on routers, better accommodating traffic patterns. Changes to the homing segments can be made without making physical changes to the network, and splitting traffic across routers and hubs protects the network against component failures.

Route Determination for ATM End Systems

There are two major functional components of an ATM switch. The first is cell-forwarding functionality, or digital cross-connect logic, which to date has received most of the attention. In general, once a virtual channel is established, ATM switches have exceptional high-speed and low latency cell-forwarding capabilities.

Figure 5. *Using an ATM Switch to Cross-Connect LAN Segments and Build Router Clusters*

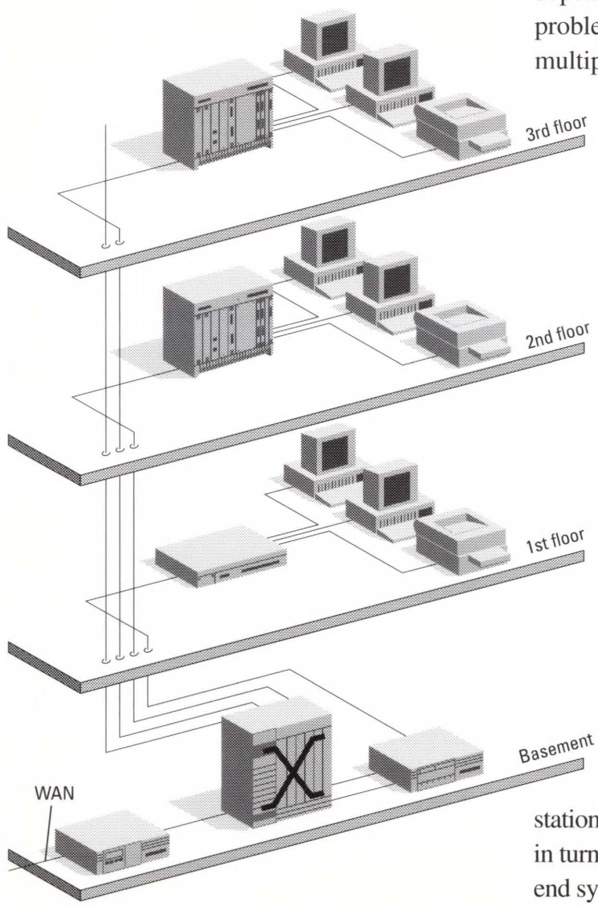

An ATM switch is connected between the floor hubs and the collapsed backbone router. Virtual connections link each router to every other router on the backbone. Connecting a cluster of routers to the ATM switch provides links to wings of a building or other buildings within a campus.

The second component is route determination functionality, which sets up the virtual channels. In the existing LAN environment, route determination is handled automatically by the routing engine in the collapsed backbone router. This level of automatic functionality is also needed in the ATM environment.

Since end systems use many different protocol stacks (such as TCP/IP, IPX, AppleTalk™, and DECnet™), the required route determination function is protocol-dependent. The easiest way to solve this problem is to leverage the 400 man-years of multiprotocol route determination logic in the switching engine of the collapsed backbone router by adding route determination for ATM along with Ethernet, Token Ring, FDDI, X.25, frame relay, and so on.

Figure 6 illustrates this simple solution. An ATM end system establishes a default virtual channel to the collapsed backbone router and transmits its route determination frames (for example, ARP frames for IP) across this virtual channel. If the destination system identified in the route determination frame is attached to the same switch (that is, it also has a default virtual channel to the router), the routing engine in the router helps set up a virtual channel between the two end systems for direct communication. Otherwise, the routing engine helps set up a virtual channel between the requesting ATM end station and the collapsed backbone router, which in turn forwards frames to and from the remote end system; the end system could be locally attached to an existing LAN or in a remote site. From an architectural perspective, ATM switching is an extension of the advise-and-consent technique.

Figure 6. *ATM Switching, an Extension of the Advise-and-Consent Technique*

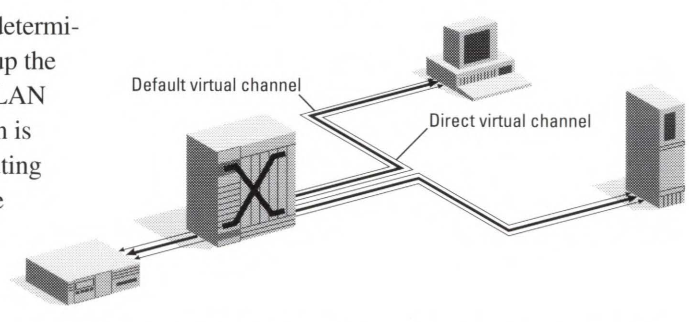

When the end system knows the route (has an established virtual channel), it forwards frames to the destination directly on the virtual channel. Otherwise, the end system requests help from the central routing engine in the collapsed backbone router, which gives advice-and-consent for the correct virtual channel to use.

Conclusion

High-Performance Scalable Networking is 3Com's strategy to meet the demands of growing networks and advanced network applications. Applied to building networks, HPSN begins with the collapsed backbone architecture and evolves through three stages:

Stage 1

- Increasing LAN segmentation
- Introducing virtual workgroups by using port grouping in the collapsed backbone router

Stage 2

- Increasing vertical infrastructure speed through high-speed LAN and ATM downlinks
- Extending the port grouping concept to virtual channels
- Implementing route caching to provide scalable routing capacity in the collapsed backbone
- Implementing local Boundary Routing or LAN switching to provide scalable network capacity on the floor

Stage 3

- Adding ATM switching by exploiting the speed and low latency of ATM
- Using the route determination of the routing engine in the collapsed backbone router

As the technology evolves, increased network capacity matches the needs of ongoing network growth, and emerging high-speed standards ensure continued control over both existing LANs and new ATM networks. Ultimately, the complex route determination logic remains centralized in the central routing engine (the NETBuilder II). Scalable performance is achieved by distributing the much simpler, high-performance advise-and-consent data-forwarding functionality to the hardware switching engines in the router, hubs, and ATM switches.

Tech Tips

Tech Tips has been developed to help you make your networks more efficient. We try to find the most meaningful technical tips from a variety of sources including 3ComFacts℠, 3Com's interactive fax service, and Ask3Com℠, 3Com's bulletin board services available on CompuServe®. Both feature technical articles, product and service information, patches, fixes, utilities, and now bug reports. Let me know how you like these tips—are they meeting your needs? I can be contacted via fax at 408-764-5001 or via the Internet at Suzanne_Dowling@3Mail.3Com.com.

Thanks to Martin Coombes, Sharon Auby, and Kathy Laymon for their contributions to Tech Tips.

Discovering Remote Subnets Using SunNet Manager

Problem:

If you are using a UNIX® platform with SunNet™ Manager to manage your network, you need a facility that allows you to easily build up your database of manageable devices on the network. This tech tip describes the SunNet utility and how it works.

Solution:

SunNet Manager has a Discover Tool that is used to build up a database of manageable devices on the network connected to each router. The Discover Tool works by pinging (calling) all possible IP addresses in its own IP subnet range. When it gets a valid response, it further interrogates the device to find out what kind of device it is. This can take a long time, up to 20 minutes, so be patient!

The Discover Tool alone will not discover devices on the other side of an IP router. When a router is discovered, SunNet Manager polls the routing table information and places icons on the diagram corresponding to each remote IP network. These icons are small "bus" lines with the IP subnet number indicated.

To discover devices on each IP network, follow these steps:

1. Click the right-hand mouse button on the bus icon to bring up the menu.

2. Under Tools, select the Discover option. Choosing Discover starts the discovery of devices on the selected network.

   You can also start the Discover process from the command line, using the following syntax:

   ```
   snm_discover -h N n.n.n.n
   ```

   where *N* = number of bits used as the device (host) address and *n.n.n.n* is the IP network number. For example:

   ```
   snm_discover -h 8 89.24.34.0
   ```

Fix for Router Forward Filters Not Functioning

Problem:

3Com has detected a problem that occurs when forward filters are used on NETBuilder® I and NETBuilder II routers running NB SW 5.1.x and 5.2.x respectively. After the filters have been implemented, they may work for a few minutes and then stop functioning.

Solution:

To fix the problem, you can upgrade your NETBuilder to software version 6.0 or later.

Note: NETBuilder I routers with only 2 megabytes should be upgraded to at least NB SW version 5.1.0.21. (At this time, version 5.1.0.23 is current.)

For more information on forward filters and documentation changes in 3Com's manual regarding the filters, refer to Ask3Com article 03190045.

Assigning an IP Address to a Boundary Router for a NETBuilder II Bridge

Problem:

Every time a NETBuilder II boundary router boots, it needs to get its IP address. There are two ways that the NETBuilder can retrieve its IP address: through the port-to-IP address mapping, or through the remote address resolution protocol (RARP) IP address translation table. Which is the correct one to use for a bridge-only environment?

Solution:

When you configure the NETBuilder as a bridge, the correct way to retrieve the IP address is through the RARP IP address translation table. Follow the steps in the manual to configure the NETBuilder II for a bridge-only environment:

1. Type: add -ip addr <*IP address*> <%*MAC address*>

   where *IP address* and *MAC address* refer to the peripheral node or boundary router.

2. Type: setd -arp rcont = rs

3. Type: setd -ip icmpr = m

Feature Article

Transcend Network Management Solution

Integrated Management Based on Connectivity Systems and SmartAgent Software

By Lionel de Maine and Bob Weder

Lionel de Maine is product line manager for network management in 3Com's NS-OPS group. Prior to joining 3Com, Lionel worked for ROLM where he held management positions in their International Telecom Division, PBX Development Group, and Advanced Call Solutions Group.

Prior to that he was a software engineer for Adaptive Networks, Inc., a manufacturer of power line–based communications equipment.

Lionel holds an M.B.A. from the London Business School, an M.S. in computer science from the University of Toronto, and a B.S. in computer science from Pennsylvania State University.

As companies' data communications requirements change and networks expand, network managers must cope with increasingly complicated multivendor environments. Traditionally, the network devices have had to be managed individually, with no way to correlate data or integrate management tools across the various devices. This puts an excessive burden on the management staff and drives up expenses.

A major time-consuming drawback of today's complex networks is the device-specific nature of LAN management systems. Without a system that can collect and pool information from a number of related devices, managers must take on the burden of collating data and turning it into useful information. The lack of an integrated management system also makes it more difficult to provide any structure for managing networks.

This article describes Transcend™ and its connectivity systems and SmartAgent™ concepts, and how these technologies can solve the challenge of managing complex networks.

Transcend's Connectivity System

To meet the challenge of effectively managing today's global data networks, 3Com is rolling out a series of innovative management products over the next two years under the Transcend name. 3Com's Transcend management solution is built on the concept of object-oriented connectivity groups and SmartAgent intelligent management software agents integrated into 3Com's adapter, hub, and router products. Transcend's object-oriented, protocol-independent architecture integrates network devices into logically related groups or objects, giving managers broad, simplified control of the network from the industry-standard management platform of their choice.

The key element in 3Com's Transcend architecture is the connectivity system, a group of related devices. Transcend takes advantage of relationships between integrated nodes to form an information management structure based on a hierarchy of intelligent SmartAgent management agents residing in the network devices. As shown in Figure 1 and Table 1 (page 14), a

Anatomy of a Device SmartAgent

A device SmartAgent includes two principal areas of functionality: (1) the management agents and interfaces necessary to support the management protocols; and (2) a generic interface to the protocols and another interface to the device-specific software that operates on the MIB data.

A SmartAgent consists of the following four components:

The Kernel. All the network management requests received by the device are routed through the SmartAgent kernel. The kernel is designed to allow various protocol stacks to access MIB items in a protocol-independent manner. It also supports MIB updates, synchronizes requests, and provides user validation using a LAN security architecture (LSA) application incorporated in the SmartAgent.

Management Agents. The interfaces to support protocols are provided by management agents. There are currently three management agents available: SNMP agents, American Standard Code for Information Interchange (ASCII) agents, and front panel agents (for devices with this feature). Agents can be added or replaced by downloading a software library.

The SNMP agent is the primary means within the SmartAgent for managing connectivity

Figure 1. *Transcend's Connectivity System Structure*

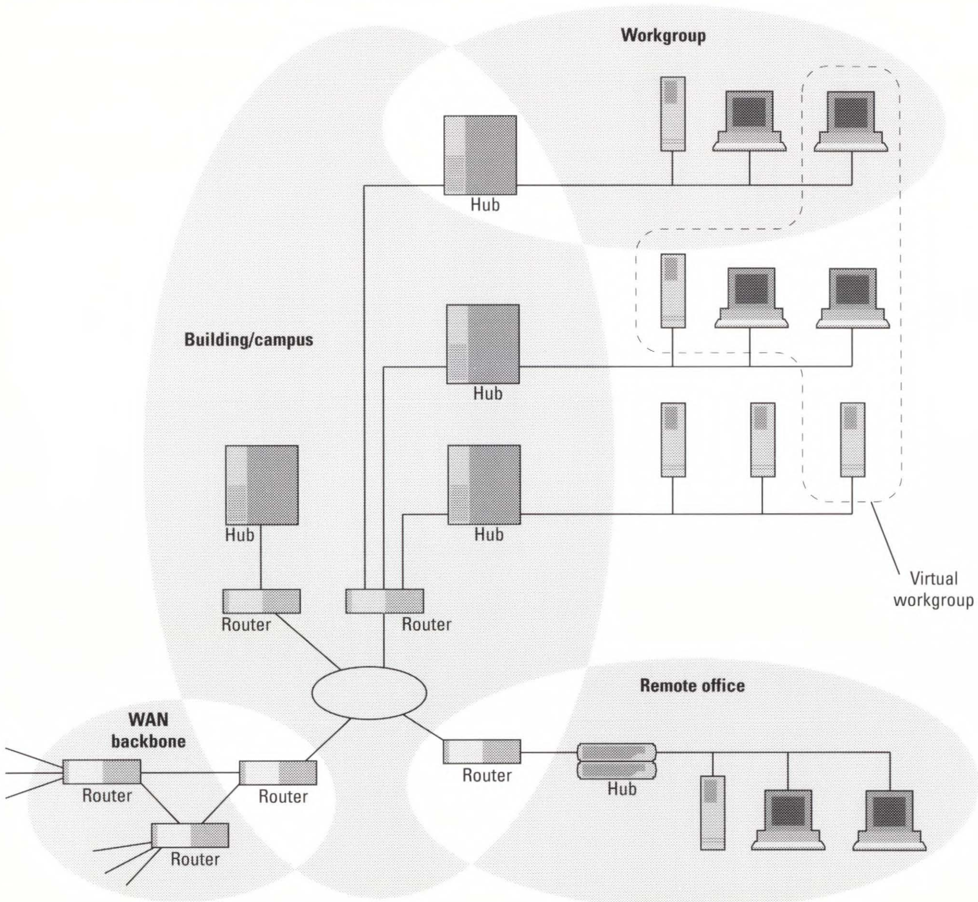

Transcend's logical connectivity system structures networks into configurable logical groups. The logical workgroup created is not dependent on physical topology.

systems and devices, and also provides the throttling mechanism that prevents error-message overload on the management station. The ASCII agent includes the Telnet protocol, a transmission control protocol (TCP) interface, and a SLIP interface for out-of-band management. It provides the manager with a user-friendly menu and form-based interface to the managed database in the SmartAgent. The ASCII agent also offers security through name identification and password protection.

System Loader. SmartAgent software may be upgraded over the network using the system loader function. This feature makes it easy for managers to add new capabilities to the network management system without having to download software at each device. It also helps preserve investment, since no hardware has to be replaced.

Gauges. The SmartAgent uses gauges to monitor statistics such as traffic levels and error counts, to measure counter rates and generate traps, and to carry out predefined actions when a threshold is exceeded. Unlike traditional network management techniques, SmartAgent gauges are autonomous from the central management station and require no over-the-network polling. Offloading statistical computation work and eliminating the need for status polling reduces the management load of the central management station.

Bob Weder is a marketing engineer for network management in 3Com's Network Management Group. Prior to joining 3Com, Bob worked for a 3Com reseller for four years, holding management positions in their network services department. He holds a B.S. in marketing from the University of Illinois.

Table 1. *Network Management Using Transcend's Connectivity Systems*

Logical Connectivity System	Types of Connections	Management Benefits
Workgroup connectivity system	Connects adapters, hubs, bridges, LAN segments, servers, and printers	• Network can be viewed in a graphical, structured way • Information can be correlated for user groups that are physically distributed among network segments or floors • Noisy devices can be located without a tedious process of elimination
Building/campus connectivity system	Connects hubs on each floor; possible collapsed backbone router and FDDI ring connections	• Generates topological maps with connectivity information on devices linked to the backbone • Booting and configuration can be accomplished from a central location • Creates and maintains asset and device inventories • Traffic analysis allows effective network design and planning
Remote office connectivity system	Connects remote routers—possibly with workgroup hubs—to a central router over a WAN link	• Remote sites can be set up and installed in a simple, plug-and-play fashion • Configuration and booting changes can be made without traveling to remote locations • Troubleshooting and problem solving can be accomplished from a central location • System can be managed out-of-band if the primary link fails
WAN backbone connectivity system	Connects geographically dispersed locations via routers attached to WAN links	• Traffic is routed according to least-cost paths • Faults occurring across the interface between the WAN carrier and the routers can be isolated and corrected • Easy links are provided to new carrier services • Accounting features track WAN line usage

connectivity system can comprise any grouping the manager wants. The ones that best match the structure of this organization are a workgroup, a building/campus, a remote office, or the wide area network (WAN) backbone. Note that the logical workgroup is not dependent on physical topology.

Central to this integrated architecture is the interrelationship between connectivity systems and SmartAgent management agents. In essence, the distributed intelligent agents turn the connectivity systems into "objects" that provide a more structured method of network control that can be shaped and managed to reflect the way an organization actually functions. Transcend management tools can then act on an entire connected system as well as on individual hubs, routers, or servers. These software tools will run on any of the major management platforms: Sun's SunNet™ Manager, HP's OpenView®, IBM's NetView®/6000, and Novell's NetWare™ management system.

The basic Transcend management structure is shown in Figure 2.

A Workgroup Connectivity System

By managing logically related groups of devices instead of individual disparate products, Transcend management software addresses the time-consuming disadvantages of device-oriented management:

- Discrete information sources that force managers to gather and collate information manually
- Device-specific management tools that require extensive product and protocol knowledge
- Multiple user interfaces that increase administrative workload and training time

To illustrate how Transcend handles these problems, we'll take the example of a workgroup. A workgroup connectivity system might be made up of any combination of PC adapters, servers, hubs, and a local bridge. Transcend automatically correlates management information for the entire workgroup, which then can be used to track trends, anticipate problems, and isolate noisy devices without going through a lengthy process of elimination.

Connectivity systems are managed with a common set of software tools that include installation, configuration, route management, administrative, and auto-topology applications. Since the tools are integrated and feature a consistent graphical interface, the network management staff does not have to spend valuable time learning diverse interfaces or manually processing information from disparate devices.

A unique workgroup connectivity systems feature is virtual node grouping. Transcend can correlate information for user groups that are physically distributed among network segments or floors, but are managed as a single virtual workgroup. As a result, nodes can be grouped together and managed according to business needs instead of their physical connections. For instance, a group might consist of everyone in various departments working on a single project. An example of a virtual workgroup configuration is shown in Figure 1, and an example of a virtual workgroup user interface is shown in Figure 3 on page 16.

SmartAgent Basics

SmartAgent is network management software capable of supporting numerous network management and communications protocols including simple network management protocol (SNMP), Telnet, and serial line internet protocol (SLIP).

SmartAgent's management and protocol interfaces are independent of specific devices. Their flexible, portable design makes it easy to integrate management functionality into an array of products. SmartAgent functionality is now in 3Com's LinkBuilder® FMS™

Figure 2. *3Com's Transcend Management System*

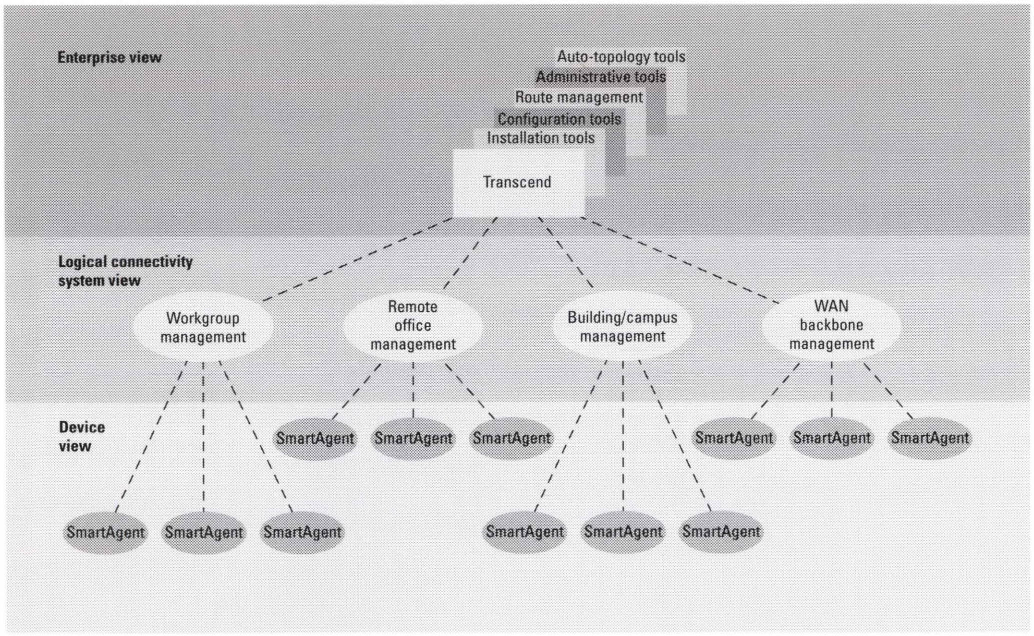

Transcend's management system structure and SmartAgent software put a high level of intelligence in the managed devices.

stackable hubs, LinkBuilder MSH™ multi-services hubs, and NETBuilder II® bridge/routers. A sample of data available from 3Com's LinkBuilder FMS SmartAgent is shown in Figure 4.

SmartAgent software puts a high level of intelligence in the managed devices. Although Transcend can also manage unintelligent devices, SmartAgent-intelligent devices provide several distinct advantages:

- They reduce the computational load on the central management workstation by intelligently interpreting traffic parameters and thresholds.
- They reduce management-related traffic on the network; central workstation device polling is eliminated because status change information is automatically forwarded.
- They reduce administrative time and labor by automating network management tasks.
- They allow devices to automatically establish a baseline representing normal operation.

SmartAgent Operation on LinkBuilder MSH

To illustrate the advantages of SmartAgent devices, let's look at error and traffic control on a LinkBuilder MSH hub.

The SmartAgent software in the LinkBuilder MSH monitors the condition of all the connected devices, the ports, and the hub chassis. If a predefined threshold is exceeded during operation, the SmartAgent automatically sends an event message, or trap, to the management station. No polling of the hub by the management station is necessary, which minimizes traffic and frees bandwidth for other uses.

If a network failure occurs that affects multiple nodes, a throttling mechanism built into the SmartAgent automatically constricts the flow of event message packets, preventing the management station from being flooded with packets.

SmartAgent's MIB object feature computes raw numerical data into useful information. For example, the LinkBuilder MSH SmartAgent calculates the percentage of overall Ethernet traffic passing through the hub, rather than merely giving a count of the Ethernet bytes received on that port.

(Continued on page 17)

Figure 3. *Transcend's Virtual Workgroup User Interface*

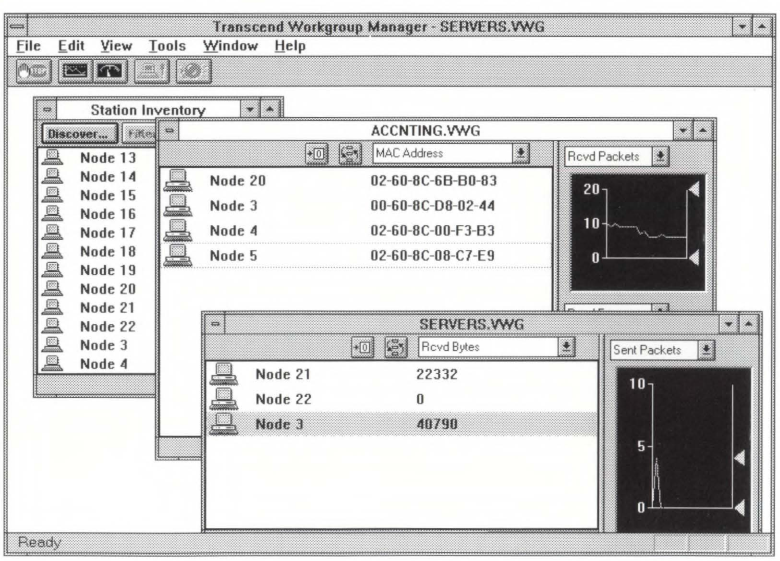

Transcend's virtual workgroup screen shows two workgroups and associated data, such as bad-packet and cyclical redundancy check information, in easy-to-read graphs.

Abbreviations and Acronyms

ASCII
American Standard Code for Information Interchange

LSA
LAN security architecture

MIB
Management information base

SLIP
Serial line internet protocol

SNMP
Simple network management protocol

Bug Reports

Now you can shortcut your network problem resolution process by getting the latest bug report with fix or workaround information any time day or night through a variety of 3Com information services:

- 3Com's CardBoard bulletin board service: 408-980-8204
  (Set your modem to up to 14,400 baud, 8 data bits, no parity, one stop bit)

- 3ComFacts interactive fax service: 408-727-7021

- Ask3Com bulletin board available through CompuServe: gothreecom (Call CompuServe at 1-800-848-8199 for an introductory membership.)

- Internet through 3Com's ftp site: ftp.3com.com (129.213.128.5)

The bug report is a listing of product software and hardware bugs, arranged by product family and organized in an easy-to-read problem-and-solution format consistent across all product lines. The bug report database is updated at a minimum of once a month.

To ensure that customers get the latest information, bug reports are published as soon as they have been verified. If you cannot find a particular bug report, it may be because there has not yet been sufficient time to determine a fix or workaround. Fixes are added as soon as they become available.

Maintenance releases and software updates are sent to all contracted service customers. If you are not a contracted service customer, contact your local 3Com sales representative to purchase the software, hardware, or firmware upgrades or releases you need.

The list that follows is a sampling of the most important bugs from the bug report database.

ADAPTERS

EtherLink III Family (3C509, 3C529, 3C579)

(Software, coax, and 10BASE-T) When an EtherLink III adapter is set with max modem speed at 38,400, the NDIS driver will have problems initializing the adapter upon a cold boot of the workstation. A PRO002E error will be displayed on the workstation at NETBIND. The solution is to set the max modem speed to 19,200 or use the latest ELNK3.DOS driver found on CardBoard in the 3C509N.EXE file.

A workaround for this problem is to immediately perform a warm boot by using CTRL-ALT-DEL keys after the initial cold boot failure.

(Hardware, coax, and 10BASE-T) If the coax or twisted pair cable is disconnected from an EtherLink III adapter installed in a NetWare 3.11 file server, the MaxPacketReceive Buffers will increment until the server crashes.

Do not disconnect the cable from the server while it is up and running.

(continued)

Bug Reports

ADAPTERS (continued)

EtherLink (3C507)

(Hardware, coax, and 10BASE-T) The EtherLink 16 is known to have hardware conflicts in EISA computers that use bus-mastering SCSI disk controllers (e.g., Adaptec AHA1542B). Revisions of the adapter prior to assy 6750-11 (coax) and 7508-04 (tp) will not work properly in conjunction with bus-mastering disk controllers. The symptom is hanging work-station of server.

To fix this problem, a hardware upgrade is needed and is available from 3Com's RMA department. A 6750-11 or 7508-04 adapter will solve this problem.

(Hardware, coax, and 10BASE-T) Assembly 6750-11 (coax) and 7508-04 (tp) adapters will not function properly in the IBM model 35sx and 40sx machines. The symptom is data corruption when copying files from a NetWare file server.

To fix this problem, upgrade the adapter to assy 6750-12 (coax) and 7508-05 (tp). This upgrade is available from 3Com's RMA department.

EtherLink/MC (3C523B)

(Hardware, coax, and 10BASE-T) IBM PS/2 models that use the 486 CPU will not function properly with the 3C523B and 3C523B-tp. Adapters prior to assy 4233-08 (coax) and 6852-10 (tp) will behave erratically. Symptoms include the inability to configure the adapter with IBM's reference disk and strange network behavior when the adapter driver is loaded (e.g, hanging workstations, network slowdowns, etc.).

To resolve this hardware conflict, contact 3Com's RMA department to upgrade the adapter to assy 4233-08 or 6852-10.

EtherLink/MC 32 (3C527B)

(Hardware) Revision of the EtherLink/MC 32 prior to assy 8209-03 is known to have a problem with transmitting an extra bit at the end of each legal Ethernet packet. Known as the "dribble bit," this can cause problems on hubs and concentrators that collect statistics on the network. These devices will report that there are many alignment errors coming from the 3C527B. The extra bit does not cause any problems from the user standpoint since the server or workstation's protocol stack will ignore the last bit after the CRC in the Ethernet packet. Everything will seem to function properly and indeed it does. Only the statistics-gathering devices will show that the extra bit is causing problems.

To fix this problem, a hardware upgrade is needed and the only way to get it is to contact 3Com's RMA department for an exchange. The boards that have this problem have assembly numbers 8209-00, 8209-01, and 8209-02.

INTERNETWORKING

6.1 Software (up to 6.1.5.02)

NETBuilder II was corrupting all received packets of sizes 1025, 1027, and 1028 bytes, only in a pure routing environment. This can happen with any routed protocol.

Fixed in 6.1.1.

OSPF slowly loses internal system messages. This leads to a failure of the NETBuilder where it may keep bridging, but there is no console response. Spanning Tree will fail.

Fixed in 6.1.5.11.

BOOTP Service: BOOTP helper doesn't work with a client on Token Ring. The NETBuilder puts the client's address in canonical format on the BOOTP reply.

Fixed in 6.1.5.09.

NETBuilder II Ethernet ports occasionally go down, then up. This only happens with 06-0085-001 Ethernet modules earlier than rev. 10.

Fixed in 6.1.5.02.

NETBuilder II Token Ring module occasionally stops transmitting.

Fixed in 6.1.1.

IPX Service: When routing IEEE to Ethernet, if the incoming packet length is greater than 60 bytes and the IPX data is of odd length, the Ethernet packet is of odd length. Some NetWare servers can't handle an odd-length packet.

Fixed in 6.1.0.03.

PLG Protocol: Packets of size > 1024 bytes do not get forwarded on the NETBuilder II when the serial line protocol is PLG.

Fixed in 6.1.5.02.

6.0 Software (up to 6.0.0.18)
LLC2 Tunneling Service could disconnect multiple sessions from the same source address when one session was disconnected from that source address.

Fixed in 6.0.0.04.

PLG Protocol Packets of size > 1024 bytes do not get forwarded on the NETBuilder II when the serial line protocol is PLG.

Fixed in 6.0.0.17.

HUBS

LB3GH version 1.1.1
Age timer not working correctly. If a user moves a device to a new ELM port, the MAC address was not forgotten in the old location. This caused misrouting of packets, lost sessions, etc.

Fixed in v1.3.2.

LB3GH version 1.3.2
When Telnetting or SNM-managing 3GH, ending a session did not release allocated memory. After approximately 18 to 20 connections/releases, no further memory available for Telnet/SNMP.

Fixed in v2.0.0.

NETWORK MANAGEMENT

LB3GH Management Software
Custom community strings ignored. When collecting FDDI ring map information for the LANplex proxy station, Viewplex would use the default strings instead of the community strings specified in the lanplex_addresses, fddi_ip_addresses, or SunNet Manager snmp.hosts file.

In 2.0.0., custom strings override the defaults for all requests.

RBCS 1.0
If a Comm Server is set to the SLP boot method, a boot appears to complete successfully, but actually hangs after the image completes loading.

The SLP daemon is sending the image file too fast. This problem is fixed with a series of (slower) SLP daemon patches, sslp1 through sslpd5, each inserting a longer delay between transmission of packets. Choose the patch you need.

RBCS 2.0
When migrating Comm Server software from NCS/AT to RBCS using bdbconv utility, Comm Servers running version 4.x code fail to boot properly from RBCS after the migration.

A new version of bdbconv was released. RBCS 2.0 patch 1.

(continued)

Bug Reports

NETWORK MANAGEMENT (continued)

ViewBuilder/UNIX/NETBuilder Management Application versions 1.0, 1.0.1, and 2.0
After reboot from tools window, further commands to NB are not allowed.

Fixed in 2.0.

TERMINAL SERVERS

CS/3000 & CS/3100 up to SW/3000-TL 6.0.0.19 and SW/3000-TLO 6.0.0.19
SIO Problems: Module lockup. On occasion an entire SIO module locks up and all ports on that specific module becomes unusable.

Ports Failing to Hang Up: Host and terminal ports don't disconnect after hangup (i.e., dropping of DTR or logout from device).

Switch Failure: Switch command fails to switch to the current session and causes the $rc (in a macro) to indicate failure.

System Lockup: Problem caused by excessive output from the user interface, resulting in continuous input on the port coming in faster than messages can be displayed. All internal system messages were used to queue up the output data, allowing no other operations to take place. Note: This condition is unlikely to occur and is recoverable. (If the input stops, all messages in the queue are eventually displayed and return to the system pool.)

System Hang (Buffer Loss): System loses track of data buffers. Condition evident to users via "Can't-no."

All problems listed above have been solved in these releases.

CS/2500 & CS/2600 up to SW/2500-TL 6.0.0.19 and SW/2500-TLO 6.0.0.19
Access Control Problems/Symptoms: Connection Validation. Rlogin and TN3270 don't validate TCP connections with the access control server if IP access control is enabled.

This problem has been solved in these releases.

Figure 4. *Example of Data Available from SmartAgent*

Transcend's LinkBuilder FMS Manager for Windows application. On the screen shown here, one window displays errors in the hub stack, the other a pie chart giving the percent of transmitted unicast, multicast, and broadcast packets as calculated by the FMS SmartAgent. With this application, a manager at a central workstation can locate hub problems and spot trends for performance tuning.

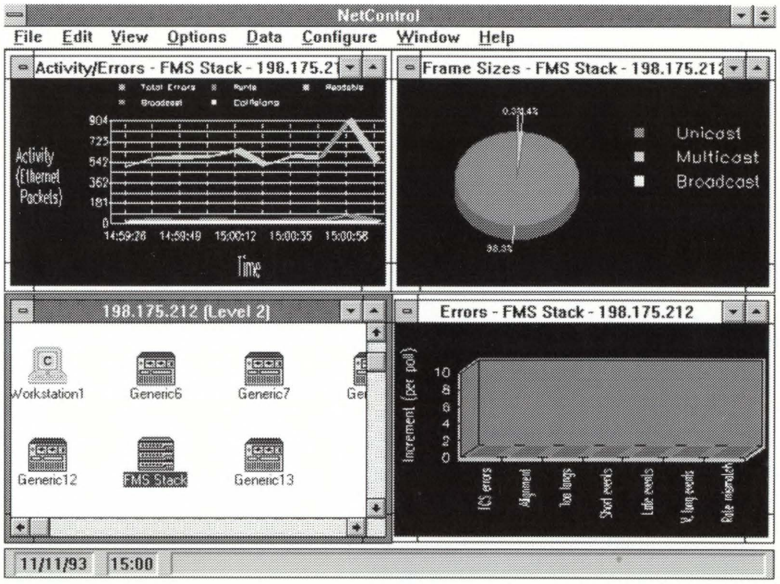

The SmartAgent software's features—interconnected gauges, autocalibration, and predefined actions—save network managers time and effort by automatically monitoring and responding to traffic thresholds.

By using any object type counter (for example, the broadcast frames received counter), a manager can set up a gauge or alarm with rising and falling thresholds that, when passed, trigger a predefined action. The SmartAgent will monitor broadcast packets and automatically partition or shut down the hub port when the rising threshold is passed, and will unpartition the port or reactivate it when the count reaches the falling threshold (representing an acceptable traffic level).

No central management station intervention or interaction is required; the SmartAgent software in the LinkBuilder MSH monitors the counter and initiates the predefined actions. As a security precaution, the SmartAgent software can also be configured to disconnect an unknown device by disabling its port. See Table 2 on page 18 for a listing of predefined actions for the LinkBuilder MSH.

SmartAgent's Autocalibration Feature

The SmartAgent's autocalibration feature makes it easier to set thresholds for a gauge or alarm. During a predefined autocalibration period, the software monitors the MIB variable and registers the high point. Autocalibration then automatically sets the rising and falling thresholds to within certain percentages above and below the high mark. (The default is 5 percent above and 20 percent below the high mark, but the values are configurable.) Figure 5 on page 19 shows an example of a typical counter over time with calculated thresholds.

The Role of SmartAgents in Connectivity Systems

The SmartAgent collects and intelligently processes information about the entire connectivity system on an ongoing basis. By designating a particular SmartAgent to collect and interpret information for the connectivity system, Transcend further extends the functionality of the agents in the individual devices.

Table 2. *SmartAgent Actions in the LinkBuilder MSH*

	Action Type	Above High Threshold	Below Low Threshold
1	No action		
2	Notify only	Send trap.	
3a	Notify and blip port	Send trap. Turn port off. Turn port on after 5 seconds.	
3b	Notify and blip module	Send trap. Turn all ports on module off. Turn ports back to original state after 5 seconds.	
4a	Notify and disable port	Send trap. Turn port off.	
4b	Notify and disable module	Send trap. Turn all ports on module off.	
5a	Notify and enable port		Send trap. Turn port on.
5b	Notify and enable card		Send trap. Turn ports back to original state.
6a	Blip port	Turn port off. Turn port on after 5 seconds.	
6b	Blip module		Turn all ports on module off. Turn ports back to original state after 5 seconds.
7a	Disable port	Turn port off.	
7b	Disable module	Turn all ports on module off.	
8a	Enable port		Turn port on.
8b	Re-enable module		Turn ports back to original state.
9	Notify and switch resilient port	Send trap. If port is the main of a resilient pair then switch to standby.	

In the case of a campus network, the SmartAgent designated for the campus connectivity system receives information from agents in the hubs and bridge/routers in the system, then correlates the information to provide a wide-ranging, cohesive view of the campus backbone.

The SmartAgent responsible for the connectivity system may be invoked automatically during daily operation, or whenever it is activated by a Transcend management tool. Information about a connectivity system obtained from its SmartAgent can be viewed as a virtual MIB for that system. The SmartAgent accepts requests for information it has already collected and responds to commands to gather additional information.

SmartAgents cooperate with each other both within and across connectivity systems, correlating and prioritizing SNMP traps and associated information from underlying device agents.

In practical terms, the Transcend architecture incorporating SmartAgents and connectivity systems gives network managers the capability to:

Figure 5. *Transcend's Autocalibration Feature*

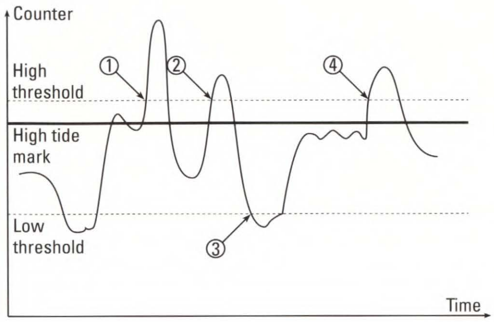

Transcend's autocalibration feature monitors a selected MIB variable and registers the high point, then automatically sets the rising and falling thresholds to within a certain percentage above and below it.

- Construct a logical connectivity system diagram providing an integrated view of the network
- Create operational baselines or thresholds and troubleshoot using the baselines
- Access error data and diagnostic routines on both ends of the network connections for automatic error isolation and correction
- Define and manage logical node groups within the connectivity systems

Conclusion

3Com's Transcend network management was designed with four major goals in mind to help solve the management complexities in today's complex, multiprotocol, geographically dispersed enterprise networks:

- To simplify and automate network management using common tools and SmartAgent functionality
- To provide integrated and scalable management of networks based on common systems
- To operate on (not replace) industry standard network management platforms
- To leverage 3Com's product breadth and expertise and provide superior network management solutions

3Com's Transcend network management simplifies the management complexity inherent in today's enterprise networks by organizing the elements of the network into logical groups of related devices. SmartAgent intelligent agents and a set of common software tools provide integrated, structured network management that will run on common industry network management platforms. Transcend's scalable, modular solution provides network managers with the tools to proactively manage network segments, improve overall network performance, and lower operating costs. □

Q+A

What Is TUBA?

Network usage on the Internet has exploded in the last decade, transforming it from an academic and research interest network into a truly global multisectored network interconnecting more than 10,000 networks, with more than 1.5 million users in over 50 countries. The original architects of the Internet did not envision this worldwide, wide-open usage, and the Internet is running out of space and routing capabilities. The fixed address size and flat routing structure of the IP datagram protocol severely limits its ability to handle the more than 2 billion Internet users predicted by the year 2000.

This Q & A describes TUBA, TCP/UDP with Bigger Addresses, one protocol solution that is actively supported with public demonstrations of implementations by 3Com, Cisco, National Institute of Standards and Technology, and Bellcore, among others, as a life extender for the IP protocol family on the Internet.

Special thanks to Cyndi Jung and Tracy Mallory for their contributions to this Q & A.

The IP-Based Internet: A Victim of Its Own Success

The current Internet IP address is a fixed, 32-bit number. Given the Internet's exponential growth projections and the fixed size limitation of the IP address, most experts agree that the number of available Internet addresses will be used up before the end of the decade.

Next to IP's address space limitation, its second major problem is the scaling of its routing protocols. Internet routing normally treats IP network numbers as a set of flat identifiers, each requiring its own routing table entry. As the number of network numbers increases, scaling problems such as memory and computational overhead for routing information, bandwidth for routing information distribution, and stability of distributed routing computations could overwhelm the current generation of routers.

Subnetting and Supernetting Can Extend IP's Useful Life

But the sky isn't quite falling—the efforts of the Internet Engineering Task Force (IETF) continue to uncover new solutions to stretch the IP support to the last inch. Two such solutions, subnetting and supernetting, have been developed to stretch the IP address space and reduce the routing overhead.

Originally, the IP address was divided into two parts, the network number part and the host part, with the value of the leading bits of the address specifying the division point between the two parts. This method is now called "classful" addressing, in contrast to the current method called "classless" addressing, where the division between the network number part and the host part is specified by another 32-bit number, a bit mask. In the bit mask, the bit value in the bit positions belonging to the network number is set to 1 and the other bit values are set to 0. Figure 1 shows the relationship of the bit mask to the IP address.

Subnetting

Subnetting has been in common use for several years, and is basically accomplished by using a bit mask that extends the network number part into the host part. The additional bits of the network number are used to provide addressing for multiple segments on the inside of the network identified by the original network number, while outside it still looks like one network and consumes only one routing table entry. Subnetting normally takes place at the end-user network and is not visible to the routers of the Internet service provider connecting the end user to the Internet.

Supernetting

Supernetting is more recently introduced and only its administrative aspects are fully supported. With supernetting, an Internet service

CardBoard—How to Reach It

CardBoard℠ is a multiple-line bulletin board system that contains the latest drivers, patches, and fixes for 3Com adapter products as well as technical articles from the 3Com technical support organization.

Example of Service:

Country	Phone Number and Modem Setting
France	*(33) (1) 69 86 69 54 Up to 9600 baud, 8 data bits, no parity, 1 stop bit*
Germany	*(49) (89) 62732-188 Up to 9600 baud, 8 data bits, no parity, 1 stop bit*
Italy	*(39) (2) 27 30 06 80 Up to 9600 baud, 8 data bits, no parity, 1 stop bit*
U.K.	*(44) (0) 442 278278 Up to 9600 baud, 8 data bits, no parity, 1 stop bit*
United States	*(1) (408) 980-8204 Up to 14400 baud, 8 data bits, no parity, 1 stop bit*

Note: Use the country code (the first number in parentheses) only when you are calling a CardBoard number outside your country.

provider is assigned a block of numbers that are related because the leading bits have the same value. In the routing tables of the routers on the public backbone, the routing information for these networks is a single entry with the accompanying bit mask that covers only the portion of the number that is the same for each in the block. The service provider is free to assign these network numbers to customers as needed, and will have to keep the network numbers as individual entries in the routing tables of their routers; but beyond the Internet service provider, these network numbers are aggregated into one entry.

Whereas subnetting is a private form of route aggregation, supernetting is a public form of route aggregation, and is the fundamental concept of classless interdomain routing, which is now fully supported by 3Com NETBuilder II products with the introduction of the border gateway protocol in the NETBuilder system software 6.2.

TUBA: An Evolutionary Protocol Based on Known Standards

The TUBA protocol attempts to solve the IP's address limitation and routing capability while continuing to support the broad base of IP users and provide the Internet with an architecture that will serve it well into the next century. The basic concept applied is the same as phase 5 of DECnet—keep the existing applications, but run them over the new network layer protocol, the connectionless network protocol (CLNP). The CLNP was developed by the International Standards Organization (ISO) for the Open Systems Interconnection (OSI) protocol family, and draws heavily from its IP and DECnet phase 4 predecessors.

Both IP and CLNP are datagram protocols, as are IPX, XNS, AppleTalk, VINES™, and DECnet. Both IP and CLNP are distinguished in this group by various scaling features, such as the ability to fragment and reassemble data packets, provide support for reporting errors, and extensibility through options.

IP and CLNP are very similar in almost every respect except for addressing. The network layer addresses used for CLNP are standard OSI network service access point (NSAP) addresses. An NSAP address is of variable length, up to the maximum of 20 bytes (Figure 1). The NSAP is structured with administrative information leading the address, followed by domain identification and a client selector. OSI addressing was designed to allow for multiple formats, as appropriate to their use. Some user groups may benefit from addressing based on public numbering schemes (ISDN, E.164, X.121, and so on), while others may choose addressing based on a national numbering scheme, independent of any underlying technology.

Unlike IP addressing, OSI network addresses are assigned to systems, not interfaces. Therefore, routers and multi-homed hosts require only one address. (There is also no conceptual equivalent of a subnet address.) Identifying hosts instead of interfaces simplifies configuration and enhances the robustness of the network. Since addresses are not bound to interfaces, if one interface on the network fails, packets can still be delivered to the same address over another interface.

Figure 1. *Comparison of IP Addresses and Bit Masks, and of IP and NSAP Addresses*

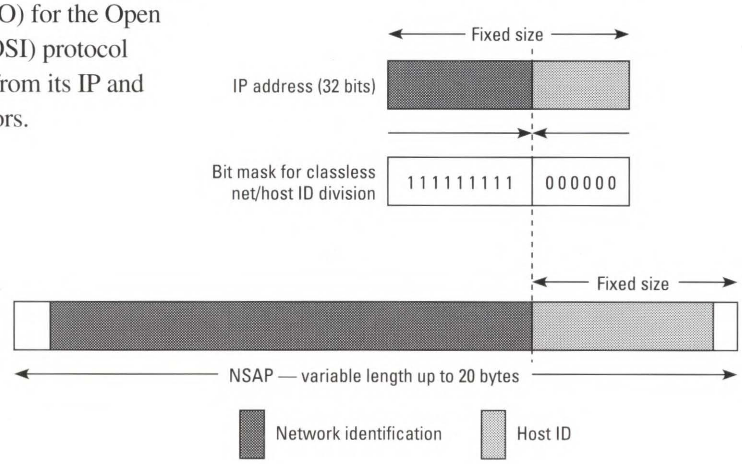

Attention 3Wizards... Get in Touch with Your 3Wizard Council!

3Wizards, the members of the 3Wizard Council are your representatives to 3Com. They are available to answer your questions and hear your views.

To contact council members, write to: 3Wizard Council, c/o Linda Webb, 3Com Corporation, P.O. Box 58145, Santa Clara,

Figure 2. *TUBA's Dual Network Layer Architecture*

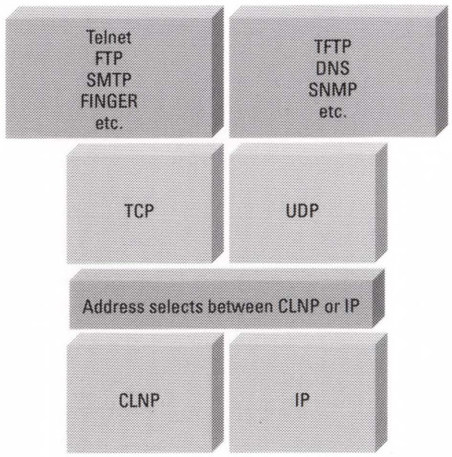

TUBA's Aggregate Routing Eliminates Routing Overhead

TUBA uses the OSI network layers full complement of routing protocols. The OSI framework consists of a global routing environment aggregated into routing domains. Each routing domain operates under a single administrative authority and is capable of efficiently handling several thousand hosts. Within each routing domain there is an aggregate of areas, each of which is a connected set of routers and hosts. Within an area, addressing and routing are flat to allow changing for host locations without having to change their addresses. An intra-area router (internal to the area) makes a simple packet-forwarding decision based on whether the destination address in the packet is inside or outside the area. If the packet is destined outside the area, the intra-area router forwards it to the nearest inter-area router.

Interdomain routing is handled by the multiprotocol interdomain routing protocol (IDRP), which is based on the IP border gateway protocol. IDRP allows many individual routes to be aggregated into a single prefix, providing several levels of routing and reducing the volume of routing data exchanged. To further simplify global routing, IDRP uses a routing domain confederation scheme, allowing sets of routing domains to be grouped as single topological entities.

Tackling the IP to TUBA Transition

TUBA's dual network layer strategy provides a relatively simple solution to the address limitations of IP, leveraging the industry's investment to date in the OSI network layer protocols IP. Figure 2 illustrates TUBA's dual structure. In the short term, IP-only hosts can use IP to communicate with other hosts. TUBA-capable hosts can use IP to communicate with IP-only hosts and CLNP to talk to other TUBA-capable hosts. Over time, IP-only hosts can be upgraded to TUBA-capable hosts if they need to communicate globally; hosts using IP strictly for local communication will not need to be upgraded.

The TUBA solution is based on existing mature standards and technologies. Many networks already carry CLNP data traffic (NSFNET, Alternet, ESnet, and NSI, for example), and most commercial routers already incorporate CLNP in their products. The large number of TUBA implementors from the U.S., Canada, and Europe demonstrating TUBA at the last several IETF meetings underscores worldwide interest in its use as a global data networking protocol. Detailed CLNP addressing and routing plans already exist to support Internet providers.

For more information on TUBA, you can join a general discussion list by sending mail to tuba-request@lanl.gov, or an implementation discussion list by sending mail to tuba@ctt.bellcore.com. TUBA-related documents are also available for anonymous FTP on merit.edu in the directory /pub/tuba-docs.

Feature Article

3Com Leads the Multivendor APPN Pack

Implementing IBM's Advanced Peer-to-Peer Networking Code on the NETBuilder II Platform

By Howard Blevins

The multivendor APPN™ demonstration at Interop in March 1993 was the latest milestone in a nine-month development effort by 3Com's IBM internetworking team. The APPN development effort continues as IBM releases more robust versions of their APPN code. This article describes the engineering issues involved in porting the first implementation of the IBM APPN code to a non-IBM platform, in this case, 3Com's NETBuilder II® platform.

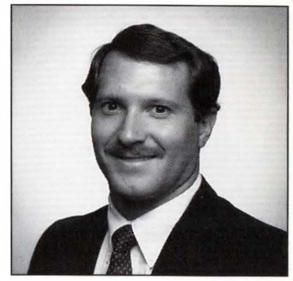

Howard Blevins is a senior engineer for APPN development in 3Com's NS-OPS division. Howard has worked for ten years in SNA networking. He received his M.S. in computer engineering from Santa Clara University and his B.S. in computer science engineering from San Jose State University.

Showcasing APPN

Visitors to the first advanced program-to-program communications/advanced peer-to-peer networking (APPC/APPN) showcase at the March 1993 Interop trade show in Washington, D.C., saw live APPN demonstrations from ten network vendors. They watched a complex Mandelbrot figure building on a screen as it was computed jointly by five other APPN systems. They could also order diagrams and lists of the names, current status, and configuration of the demonstration network and its systems.

3Com, DCL, and IBM provided the APPN network nodes, including a remote virtual telecommunications access method (VTAM) network node back at IBM's site in Raleigh, North Carolina. DCA and Insession demonstrated end nodes, while DCL, Eicon, Netlink, NSA, Novell, and SunConnect demonstrated low-entry networking (LEN) nodes. The multivendor showcase marked a significant milestone for 3Com's leading participation in software development supporting APPN code (see Figure 1).

The Year Before

Eight months earlier, in August 1992, the 3Com IBM internetworking team, including the writer, were halfway through the process of porting the APPN network node code to the NETBuilder II platform. The programming environment and the LAN 802.2 data link interface had been ported. The next step was to adapt the APPN code to run on 3Com's NETBuilder® bridge/router.

Figure 1. *March 1993 Interop APPN/APPC Showcase Configuration*

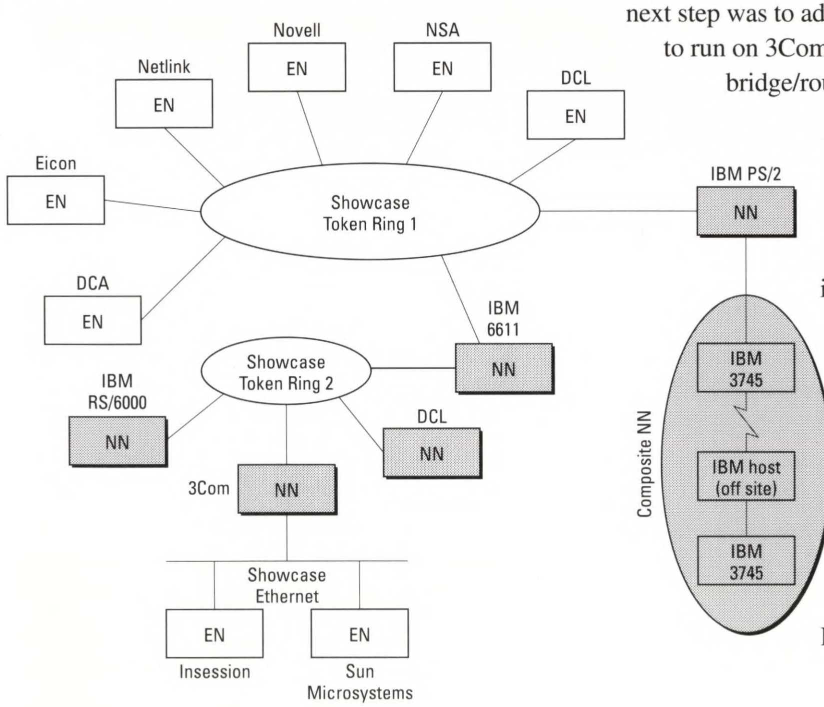

The team's internal goal was an ambitious one: to complete the port in time to demonstrate at the October 1992 Interop show in San Francisco. The challenge was formidable: 120,000 lines of C code in more than 1000 files. Long hours, sound

Abbreviations and Acronyms

APPC
Advanced program-to-program communications.

APPN
Advanced peer-to-peer networking.

DLCX
Data link control exchange.

EGPE
Extended generalized programming environment.

LEN
Low-entry networking; early APPN node type. Must be manually registered and manually configured with all potential destinations.

LLC2
Logical link control–type 2 IEEE standard.

LU
Logical unit. SNA/APPN term for software component that represents a network user—either a person or an application.

NetID
Network identifier.

SAP
Service access point.

SNA
Systems Network Architecture. IBM strategic networking architecture. Subarea SNA is precursor to APPC and APPN, the new SNA.

TDU
Topology database update.

VTAM
Virtual telecommunications access method. IBM mainframe implementation of SNA.

XID
Exchange identifier. SNA and APPN process to exchange capabilities and negotiate session characteristics before session activation.

source code from IBM, frequent IBM consultations, good tools, and the 3Com team's strong grounding in IBM internetworking combined to achieve the very first multivendor network node–to–network node demonstration. It was actually an unscheduled event: 3Com had received the APPN network node code from IBM less than three months before the show and few believed a functional port could be accomplished in time.

But two weeks before the show, IBM and 3Com ordered Token Ring cables to be pulled under the trade show floor between their booths. As Interop Fall 1992 opened, Mandelbrot floating-point fractals were being calculated between the two booths.

When It All Began

The seeds of the project were planted four years earlier, when 3Com acquired Communications Solutions, Inc.; with that merger, the core IBM internetworking team joined 3Com. The engineers developed and maintained 3Com's relationship with IBM's APPN architecture team. The ongoing working relationship led to IBM's selection of 3Com, in March 1992, to be the first vendor to port the APPN network node code to a non-IBM platform.

In April and May 1992, the 3Com-IBM internetworking development team scoped out the porting process. Starting in June, the team traveled several times to IBM in Raleigh, North Carolina, for training, particularly on the main processes, internal process flow, and data structures. Then 3Com received the beta code and the real work began. (See Table 1 for the chronology of 3Com's APPN participation.)

There were several distinct steps in the process of porting IBM's network node code, originally designed for the IBM 6611 router, to 3Com's NETBuilder platform:

- Map extended generalized programming environment (EGPE) onto 3Com NETBuilder kernel calls
- Port data link control exchange (DLCX) interface to 3Com LLC2 for LANs
- Port processes up to XID negotiation in software environment
- Port processes from bind negotiation on three-node hardware environment

Porting the Programming Environment

The first step was porting the EGPE. EGPE provides generic process calls as an interface between APPN and the platform's operating system. This IBM generalized interface enhances portability to non-IBM platforms.

For this project, porting EGPE consisted primarily of mapping EGPE to 3Com NETBuilder kernel calls. For example, to create a process, EGPE says Pcreate, while the 3Com call is pro_create. However, this port involved much more than simple function mapping. A single call in EGPE might map to multiple calls for the NETBuilder, or vice versa. In addition, some calls had different numbers of arguments and options.

The EGPE also involved several timer routines and buffer management procedures. Porting EGPE was relatively straightforward. The process was aided by IBM's APPN EGPE Monitor, a sophisticated command interpreter. The 3Com team used the software as a debugging tool to look into process control blocks, observe process status, and check APPN usage.

Porting the Data Link Interface

The next step was porting IBM's DLCX interface to 3Com's LLC2 stack. DLCX allows the APPN code to talk to the data link drivers. All of IBM's code was written to generic UNIX® driver calls, such as I/O control [ioctl()], and none of these applied to the NETBuilder.

DLCX provides the means to configure a port to use APPN. It creates a reader and a writer when a port is activated. The reader passes to APPN any data the node receives from LLC2 containing an APPN service access point (SAP). The writer, in contrast, passes to LLC2 any data from APPN destined to the APPN

Table 1. *Timeline for 3Com APPN Support*

1988–1992	Develop and maintain relationship with IBM development
March 1992	Announcement of OEM relationship for APPN network node
June 1992	Training at IBM Networking Systems in Raleigh, NC
July 1992	Received beta APPN OEM code
October 1992	First multivendor APPN network node demo at Interop Fall, San Francisco, CA
February 1993	Two hot stagings in Raleigh, NC, before Interop Spring
March 1993	APPN Showcase at Interop Spring in Washington, DC
April 1993	APPN network node licensed code and specifications available
May 1993	Repeat port process with official code
August 1993	3Com hosts Interop Fall hot staging in Santa Clara, CA
August 1993	APPN multi-booth APPN demo at Interop Fall, San Francisco, CA
September 1993	Host staging in Raleigh, NC, before Interop Europe
October 1993	APPN Showcase at Interop Europe in Paris, France
January 1994	3Com to ship APPN Network Node for NETBuilder II Ver. 1.0
May 1994	Scheduled APPN Showcase at Interop Spring, Las Vegas, NV

SAP. A separate reader and writer must be developed for every data link type. Fortunately, Token Ring, Ethernet, and FDDI all use the same 802.2 LLC2 data link control interface.

Translating from a Multiple to a Single Code Space

After developing the interface to the programming environment and the data link control, the process of porting the APPN code to a real node began. The reliability and structure of IBM's code helped to make the port a fairly straightforward task.

But as in any porting process, the 3Com team uncovered bugs along the way. They reported the bugs to the IBM APPN development engineers on a regular basis, and the IBM team quickly worked to debug them. If necessary, three or four developers familiar with the different aspects of a discovered bug teamed up for a conference call with 3Com to resolve the problem.

A major obstacle to the porting process lay in a difference between the IBM 6611, from which the code was derived, and the 3Com NETBuilder. The 6611 is an AIX® system, IBM's version of UNIX, and AIX contains a UNIX kernel that runs the multiple processes in separate code spaces. Therefore, since no code space is common, several processes can use the same variable names for different purposes.

The NETBuilder, on the other hand, is a multiprocess system that runs the APPN code in a single code space. Therefore, no variable name duplication could be allowed; all name conflicts had to be resolved.

The 3Com team handled this issue by affixing a two- or three-letter acronym to the beginning of each variable when a collision was found. The acronym was an abbreviation for the component name where the variable was used. For example, names related to the topology and routing services were prefixed with TRS_ and names related to the session connector manager were prefixed with SCM_. This process involved finding and modifying each occurrence of these variables throughout the code.

Reducing Variable Search Time with the Identifier Database

The task of finding and if necessary modifying each variable in all files in which it occurred was daunting. Developers usually use the UNIX *grep* (global recognition of expected pattern) command to find all instances of a particular string in a file or set of files. But to search the APPN code, *grep* would thoroughly examine each of 1000 files to find the few files in which the string occurred. Each search might take up to an hour.

Fortunately, the 3Com team had for years been using a freeware tool available on the Internet, the identifier database (ID database). The ID database creates a database of all identifiers—including function names, structure names, constants, variables, and defined items—for a set of files, as shown in Table 2. This database notes all files where each identifier is found. Other commands can use this database to significantly reduce search times.

The make identifier database (*mkid*) command was used initially to build the ID database and to rebuild it after changes were made. To *grep* a variable, the 3Com team used grep for identifier (*gid*) on this ID database instead of on all 1000 files. The *gid* output listed the path and the line of code for each occurrence of that string. This list could be used to note where duplications existed.

Once the duplications were noted, edit identifier (*eid*) was used to edit every occurrence of duplicated variables. Determining the appropriate prefix to use was relatively straightforward from the pathname that noted the directory containing the file.

Debugging the XID Negotiation Process

After all the variable names were adapted to work in the NETBuilder II environment, the 3Com team turned to debugging the exchange identifier (XID) negotiation process. In APPN, the XID process manages all communications about capabilities and preferences between two nodes before any sessions are activated.

To minimize and isolate problems, this portion of the port was done completely in software. Using 3Com proprietary component workbench software on a Sun SPARC® workstation, the 3Com team developed a loopback driver to simulate XID negotiation on a single platform (Figure 2).

Porting to the NETBuilder II Hardware Platform

After XID negotiation, two APPN nodes bind a session between their control points (CP–CP session). The bind process would have been too complex to debug using a loopback driver, so the 3Com team moved on to port directly to the NETBuilder II platform.

Three computers were used for the initial port, as shown in Figure 3. Two OS/2® systems ran IBM Communication Manager. One system was configured as an end node, to test the NETBuilder as a network node server. The other was configured as a network node, to test topology updates and searches with the

Table 2. *ID Database Identifiers*

asm_mainlin	pc/{xxxams00,xxxpcasm}.c
cs_main	cs/xxxcss00.c nof/xxxnof01.c
dlcx_main	nof/xxxnof01.c dlcx/cswdx000.c
dll_main	ps/xxxpss00.c
ds_main	nof/xxxnof01.c appnds/newdss00.c
hs_main	hs/xxxhss00.c nof/xxxnof01.c
ms_main	ms/xxxmss00.c nof/xxxnof01.c
pc_main	nof/xxxnof01.c pc/xxxpcasm.c
pc_mainline	pc/{xxxpcasm,xxxpcs00}.c
ps_main	ps/xxxpss00.c rm/{xxxrmp35,xxxrmp44}.c
rm_main	nof/xxxnof01.c rm/xxxrms00.c
scm_main	nof/xxxnocsp.c scm/xxxims00.c
sm_main	nof/xxxnolud.c sm/xxxsms00.c
ss_main	nof/xxxnof01.c ss/xxxsss00.c
ssa_main	nof/xxxnof01.c ssa/cswsas00.c
tr_reader_main	dlcx/{cswdxsds,cswdxtrw,cswdxtrr}.c
tr_writer_main	dlcx/{cswdxsds,cswdxtrw}.c
trs_main	nof/xxxnof01.c trs/xxxtos00.c
tun_reader_main	dlcx/{cswdxtun,cswdxtur,cswdxtuw}.c
tun_writer_main	dlcx/{cswdxtun,cswdxtuw}.c

Sample ID output shows some of the names containing the word "main" and the main entry points for the various processes in APPN.

Figure 2. *Development Platform for XID Negotiation*

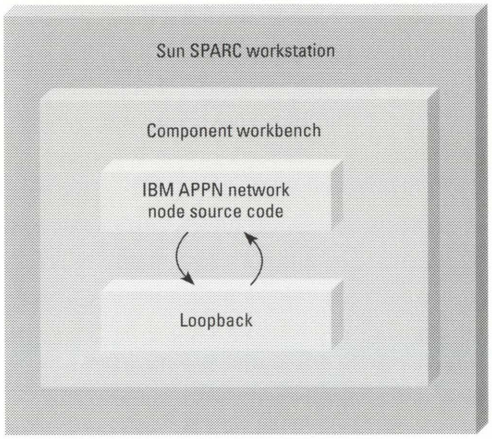

NETBuilder. Both the OS/2 end node and network node contained applications to test logical unit–to–logical unit (LU–LU) session establishment managed by the NETBuilder network node.

Handling Memory Challenges

The 3Com engineers discovered three memory challenges in porting APPN to a multiprotocol router. First, 120,000 lines of code take up a lot of room, more than any other single protocol on the NETBuilder.

Second, in the current APPN network node code, the buffer manager preallocates memory for the entire pacing window. For example, a pacing window of 7 with a request/response unit size of 4000 requires 28,000-worth of buffers. This 28,000 amount is allocated to that session, so it is not available for any other session, even if no data is flowing.

The third challenge is related to a difference between the 6611 and NETBuilder platforms. The number of end nodes a network node can support depends on the memory available. Since the 6611 AIX has virtual memory, its APPN can operate as if it has 25 Mbytes for itself, without taking up 25 Mbytes of actual memory. Actual memory is used only as needed. The 3Com NETBuilder and most other routers do not use virtual memory because it can slow down a router operating in real time. Therefore, the NETBuilder allocates actual memory to APPN, which results in an optimal implementation of APPN, focusing on performance rather than supporting very high numbers of sessions.

As APPN develops, memory usage will improve, using statistical buffering methods to increase the number of sessions supported through buffer allocation based on actual recent usage.

Instilling Interprotocol Fairness

Because of the differences in the way the IBM 6611 and other multiprotocol routers handle CPU processing, the 3Com engineering team developed a fairness equation to allow APPN to share NETBuilder resources fairly with other protocols. The 6611 AIX system handles multiple protocols with the UNIX capacity for time slicing. UNIX performs context switches, switching between many simultaneous processes in each time slice.

Most multiprotocol routers, however, are not built for preemptive multiprocessing. Instead, a single process holds the CPU until it is completed. This is not a problem with most protocols, which are designed to use the CPU very briefly for each process. However, because the 6611-based APPN code expects the operating system to control usage of the CPU, it acts aggressively in a nonpreemptive environment; it assumes it can use the CPU until told to release it.

Figure 3. *3Com APPN Porting Configuration*

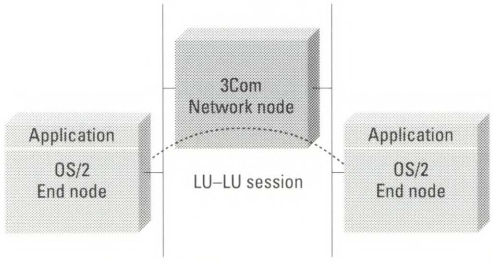

A logical unit–to–logical unit session (OS/2 network node to OS/2 end node through 3Com network node)

After extensive testing, 3Com engineers developed an equation to grant a fair number of CPU cycles to APPN before it must give up the CPU in favor of another process.

First Multivendor APPN Demo

From in-house development, the next step was to test the code in a true multivendor environment. Only 3Com and IBM participated in the October 1992 Interop APPN demonstration, but the Interop show the following spring featured ten vendors demonstrating APPN connectivity. This multivendor demonstration provided 3Com, as well as the other vendors, with a golden opportunity to test their code in a rigorous environment.

In preparation for the March 1993 showcase, IBM hosted a hot staging in February in Raleigh, with the author as the APPN/APPC showcase technical chairperson. Since it was the first multivendor APPN demonstration, several interoperability problems were uncovered.

One problem was caused by two vendors, Insession and IBM, who interpreted the APPN documentation differently. IBM's end node documentation states that an end node registers itself and its resources to the local network node using local LU names. Insession interpreted the documentation to mean that resources can be registered with non–fully qualified LU names. But the IBM code expected the LU names to be fully qualified names and did not register them properly. Consequently, searches for those resources failed because they were improperly registered. Both interpretations were technically correct, but it took an interoperability demonstration to discover the discrepancy. As a result, Insession modified their code to support fully qualified LU names and IBM modified the code to support non–fully qualified LU names.

Not all the bugs were resolved before Interop, however. The remote-VTAM-host connection to Raleigh caused some problems that brought the showcase network performance to a halt several times during the show. The mainframe, which was running new code, generated topology database updates (TDUs) containing a new, undocumented field unknown to network node implementors 3Com, DCL, and IBM. These network nodes dropped this unrecognized variable before propagating their TDUs to the other network nodes. When the VTAM host received the altered TDUs, it noted that information was incomplete and interpreted it as an error. The host therefore resent its TDU to every network node, which again propagated it without the new variable. Due to APPN's topology update efficiency, this process repeated itself quite frequently—hundreds of times per second, congesting the network so that little or no data could be sent. As a temporary fix, IBM disconnected some nodes connected to VTAM back in Raleigh so that VTAM would not send TDUs for those nodes to the showcase. IBM later released updated APPN code that fixed the bug.

Conclusion

Multivendor APPN is here today and 3Com is leading the pack. 3Com was the first network vendor to port IBM's APPN network node code, and 3Com demonstrated its accomplishment at the Interop trade show in October 1992, communicating smoothly with IBM APPN nodes.

The porting process involved several steps: mapping the interface to the operating system, EGPE, and the interface to the data link, DLCX, to the NETBuilder platform; debugging the XID negotiation; and implementing the code and debugging the session initiation process on a three-node network using the actual NETBuilder platform.

3Com has continued in a leadership role, participating in interoperability demonstrations with fifteen other vendors at three subsequent Interop shows. 3Com's selection as IBM's partner in porting APPN network node code reflects the company's stature and expertise in IBM internetworking, including subarea Systems Network Architecture (SNA), Token Ring, and APPC. The company's relationship with IBM ensures that 3Com technology will track closely with IBM as that company evolves its SNA and APPN strategy.

Product Update

New Products at a Glance

Here's a quick view of new 3Com products that have shipped in the last nine months. Some of these products are featured in 3TECH *articles. Hope this helps you in your product evaluations.*

Adapters

FDDILink™-UTP adapter

A 32-bit EISA adapter for FDDI-over-UTP (CDDI).

- Complete implementation of the ANSI X3t9.5 TP-PMD 100-Mbps standard
- Features same "slice" technology and drivers as current FDDILink adapter
- Low-cost CDDI solution simplifies FDDI cabling installation

TokenLink® III EISA adapter

- 100-percent IBM-compatible with a money-back guarantee
- RPL utility for easy driver downloading (which is unique among Token Ring adapters)
- On-board DB-9 and RJ-45 ports
- Fully SNMP-ready through Transcend TokenLink SmartAgent software

Bridge/Routers and Software

NETBuilder II Enhancements

An award-winning multiprotocol, high-bandwidth bridge/router that supports multiple Token Ring, Ethernet, FDDI, and WAN connections.

- HSSI high-speed WAN module supports 52-Mbps full-duplex line speeds including E3, T3, and fractional line rates; provides support for future ATM implementation
- Ethernet two-port fiber module doubles NETBuilder II's Ethernet LAN port density; built-in transceiver provides direct FOIRL-to-router connection
- (Optional upgrade) Hot-swappable, load-sharing dual power supply provides fault tolerance and uninterrupted reliability for critical networks
- (Optional upgrade) Flash memory drive solid-state storage and boot device based on PCMCIA standard; includes 4-megabyte Flash RAM card

NETBuilder II 6.2 Software

Bridge/router software that supports multiprotocol WAN connectivity and Boundary Routing architecture. Runs on the entire NETBuilder family.

- Event scheduler schedules dial-up WAN connections for recurring or singular events; secures remote offices from unauthorized after-hours use; takes advantage of cheapest tariffs based on time of day
- Non–full-mesh frame relay supports dynamic routing across entire frame relay network, avoiding cost of full-mesh PVC connectivity
- Enhanced SMDS protocol support: RIP (IP, IPX, XNS, VINES); AppleTalk, DECnet, OSI, and transparent bridging support
- In-band file transfer allows download of software images and configuration files from central to remote LANs to disk or Flash floppy (NETBuilder II only)
- Enhanced filtering service allows administrator to build a table of users that can be referenced by a single group name for filtering purposes
- Expanded MIB support: Ethernet, Token Ring, frame relay, mnemonic filter, prop. bridge extension, prop. system, prop. IPX, and RMON; updated IP

NETBuilder Token Ring (TR) Remote Access

A remote Token Ring LAN/WAN router that offers the flexibility of Boundary Routing technology or conventional routing for small to midsized remote sites.

- A full routing device with one Token Ring LAN and one WAN connection

(continued)

3Com Quick Technical Resource List

Product Information
800-638-3266; outside the U.S. and Canada, call your local 3Com sales office

Technical Support
800-876-3266, press option 1; outside the U.S. and Canada, call your local 3Com sales office

Service Contract Help Line
800-876-3266, press option 3; outside the U.S. and Canada, call your local 3Com sales office

Training and Independent Study Courses
800-876-3266, press option 7
Fax: 408-764-7290
Outside the U.S. and Canada, call your local 3Com sales office

Ask3Com Information Service on CompuServe
800-848-8199; outside the U.S. and Canada, call the nearest CompuServe office

3Wizard℠ Program
408-764-6764, or contact Education Services at 800-876-3266, press option 7

3TECH *Journal*
800-638-3266, press option 4, ext. 5898; outside the U.S. and Canada, call 408-764-5898

Automated Fax Service

3Com's interactive fax service, 3ComFacts, provides technical information on 3Com products 24 hours a day, seven days a week. To access this service, dial 408-727-7021 from anywhere in the world using your touch tone telephone. In Europe call (44) 442 278279.

Free local access is available within the following countries using the numbers below:

France
05 90 81 58

Germany
0130 81 8063

Italy
1678 99085

The Netherlands
06 0228049

Sweden
020 792954

United Kingdom
0800 626403

United States
1-408-727-7021

Bridge/Routers and Software (continued)

- Conventional bridging and routing decisions are made at the remote site
- Supports industry-standard protocols
- Upgradable to NETBuilder Remote with just a software change

NETBuilder Token Ring (TR) Remote Control

A remote Token Ring LAN/WAN router that implements Boundary Routing system architecture for smaller remote sites.

- Routing decisions are made at central router
- Eliminates remote administration
- Requires central NETBuilder II running Boundary Routing software version 6.0 or higher
- Upgradable to NETBuilder Token Ring Remote Access or NETBuilder Remote with just a software change

Hubs

LinkBuilder FDDI

Cost-effective, four-slot FDDI workgroup hub supports one management module and up to three port adapters: four-port fiber-optic, six-port shielded twisted-pair, or eight-port unshielded twisted-pair (category 5) modules. Mix cabling modules to meet current and future needs.

- Dual, single, or dual-homed attachment to an FDDI ring, or stand-alone hub (null attach configuration)
- Management via SMT 7.2, Telnet, local console, or modem attached to management module's serial port
- Compliant with ANSI specifications for FDDI over fiber and the draft twisted pair–physical medium dependent standard for unshielded twisted-pair capability
- Flash EPROM provides easy software upgrades via the network, with no need to replace a ROM chip

LinkBuilder Flexible Media Stack (FMS) Fiber Hub

Six-port fiber hub, fully 10BASE-FL compliant; can be integrated with coaxial and twisted-pair FMS hubs in a single stack.

- Has an AUI port and transceiver module for coax or fiber backbone connection
- Full SNMP management to MIB II and 3Com extensions via a single module per stack, which can be added as needed and installed by users
- Can be upgraded to new management protocols via a software download

LinkBuilder Multi-Services Hub (MSH)

An eleven-slot, chassis-based hub that uses a high-speed passive independent backplane to combine Ethernet and Token Ring LANs in one managed network.

- Combines up to three Ethernet segments and five Token Ring workgroups, managed via a single management module
- Has a passive, high-speed, slot-independent backplane supporting up to 1.7 Gbps
- Distributed management architecture and SmartAgent software for cost-effective and future-proofed SNMP management
- Phase-locked loop retiming and RingBuilder™ architecture ease configuration and increase reliability on Token Ring networks
- Distributed Repeater Architecture and LAN security architecture for security
- Designed to support FDDI, ATM, and other high-speed technologies in the future

LinkBuilder Third-Generation Hub (3GH) Ethernet Switch Module

Switching or bridging mode modules for 3Com's LinkBuilder 3GH hub; they replace the Ethernet LAN Module (ELM).

- Ethernet Switch Modules support switching or bridging modes within the same network; supports both twisted-pair and fiber
- Ethernet switching mode eliminates packet flooding caused by address learning and aging
- Bridging mode configures hub as fully compliant IEEE 802.1d bridge for LAN-to-LAN interconnection

LinkBuilder Third-Generation Hub (3GH) CDDI Module

Twelve-port FDDI-over-UTP (CDDI) module; built to comply with the most recent ANSI TP-SMD draft standard.

- Ideal solution for a collapsed backbone connection to a cluster of servers
- Complements 3Com's line of FDDI-over-UTP adapters
- Recent LANQuest test of FDDI shows a tenfold increase in data throughput over Ethernet

LinkBuilder Token Ring (TR)

A stackable, manageable, 16-port Token Ring hub for shielded and unshielded twisted-pair.

- Can be used as an unmanaged hub or as an expansion unit for one of the managed (TRi) models
- Supports a maximum of 256 users per ring configured as 16 stackable hubs
- Allows a maximum of fifteen stackable hubs to be managed from one managed stackable hub (TRi)

LinkBuilder TRi

A stackable, managed, 16-port Token Ring hub for shielded and unshielded twisted-pair.

- Model 1—a managed hub that supports six hubs in a stack with SNMP and basic RMON management
- Model 2—a managed hub that supports 16 hubs in a stack with SNMP and full RMON management

Network Management

Transcend TokenLink SmartAgent

SNMP remote management for all TokenLink III and IBM adapters.

- Provides PC hardware and software information, user information, traffic and error statistics, and remote adapter disabling
- Partial manageability (user information and error statistics only) for all IBM Token Ring adapters
- Has an impact on node memory of less than 4.5 KB

Transcend LinkBuilder Token Ring Manager

Graphical SNMP-compliant network management software for LinkBuilder TR and LinkBuilder Focus™ Series hubs.

- Comprehensive network management for UNIX (NetView 6000, SunNet Manager, and HP OpenView) and OS/2-based systems
- Supports management and monitoring of 3Com and third-party SNMP devices
- Supports SNMP full RMON MIB on both Token Ring and Ethernet LANs

Transcend LinkBuilder FMS Manager for Windows

Network management software for Microsoft Windows™ 3.1 platform for small networks without full-time administration.

- Graphical point-and-click operation; hierarchical network map with realistic view of FMS stack and imported background pictures
- Supports SNMP over IP and IPX; features auto-discovery of all SNMP devices
- User passwords with access levels prevent unauthorized access
- Color-coded icons, audible trap warnings, and ping support provide quick and easy troubleshooting
- Runs on standard NDIS packet drivers on EtherLink® adapters

Transcend NETBuilder Manager 2.0

Network management software displays real-time view of NETBuilder bridge/router products, up/down interface status, and configuration information.

- Graphically configures bridge and routing parameters: bridging, IP routing, IPX routing, AppleTalk routing, packet filtering, and FDDI, Token Ring, and high-speed serial (HSS) interface media
- Easy-to-read graphical meters make it easy to spot trends and anticipate problems before they occur
- Requires NETBuilder 6.2 software

3Com Professional Services

3Com provides professional services directly and in partnership with resellers worldwide to assist customers with large or complex network projects. These services include:

- ***Network integration services***
  Services include needs analysis, design, project management, installation, cable plant certification, and start-up training.
- ***Consulting services***
  Assistance performing scheduled tasks: planning, custom training, and performance auditing.
- ***Project services***
  Well-defined projects that must be completed in a limited time: protocol migration, cable plant surveying and mapping, and WAN or LAN/campus network design.
- ***Installation services***
  Placement, configuration, verification, and connection of 3Com products to your network.
- ***Technical training***
  Independent study and classroom courses on industry-standard technologies and 3Com products.

For more information about 3Com's professional services, contact your local 3Com sales office. See the inside back cover for a list of 3Com U.S./Canada and worldwide sales offices.

Training Information

To order a course catalog or for information on scheduling, course content, or prerequisites, call 800-876-3Com, press option 7. Outside the United States and Canada, call your local 3Com sales office.

3Com Training and Information Resources

Interactive Fax Service Now Covers All Products

3Com's interactive fax service, 3ComFacts, now offers technical support information on all 3Com product lines 24 hours a day, seven days a week. Due to the success of the CardFacts℠ adapter fax service, 3Com has introduced NetFacts℠, an information service for network systems products: LinkBuilder hubs, NETBuilder routers, CS terminal servers, and supporting product software.

New documents are added to 3ComFacts on the 15th day of every month. The information service will expand in the future to include technical bulletins, application notes, and performance comparison data.

You can reach 3ComFacts by calling 408-727-7021 in North America or 44-442-278279 in Europe. Once inside the system, you can choose between the two product line services:

- CardFacts has adapter installation diagrams, configuration drawings, troubleshooting instructions, data sheets, technical articles, tech tips—and just introduced, bug reports. Select document #9999 for a complete index of the CardFacts contents.
- NetFacts has data sheets, tech tips, technical articles, and bug reports on hubs, bridges, routers, terminal servers, and software products (see sample below). Select document #8888 for a complete index of the NetFacts contents.

In addition to 3ComFacts, 3Com offers technical information through other service venues: SOS for Networks, a multivendor support CD-ROM database on Ziff Communication's Computer Library; and the Ask3Com bulletin board service through CompuServe.

Here is a small sampling from the NetFacts index:

- LinkBuilder FMS/10BTi: Enterprise trap 25 and how to disable it (#6233)
- Adjusting the LinkBuilder 3GH power supply for 110 or 220 volts (#2198)
- Using the resilience function on an ECS management card (#6126)
- Terminal server line card: SNMP management and remote console facility (#6175)
- Cabling a terminal to CS/3X00, LinkBuilder 3GH, or NETBuilder II (#2226)
- Duplicating comm server parameter information to new comm servers (#2229)
- StarTek™ 828AT-5 Intelligent Multi-station Access Unit (datasheet) (#7266)
- NETBuilder responds incorrectly to SNMP request of IP routing table info (#2238)
- Creating aliases for NETBuilder commands (#2212)
- NETBuilder macro to verify stable routes through an IP network (#2164)
- List of agencies that have approved NETBuilder II (#2178)
- Product tips on FMS and 10BTi (#6248)
- Terminal server line card: important configuration advice (#6243)
- Support requirements for ViewBuilder UNIX applications (#2275)
- ISOLAN repeaters: the logical choice for linking LAN segments (datasheet) (#6251)

3Com U.S./Canada Sales Offices

3Com HEADQUARTERS
Santa Clara
Phone: 800-NET-3Com
or 408-764-5000
Fax: 408-764-5001

3Com U.S.A.
California
Costa Mesa
Phone: 714-432-6588
Fax: 714-432-6574

Los Angeles (Downtown)
Phone: 213-892-6343
Fax: 213-892-2203

Los Angeles
(LAX Airport)
Phone: 310-348-8110
Fax: 310-348-8167

Pleasanton
Phone: 510-847-2040
Fax: 510-463-1560

San Diego
Phone: 619-546-4476
Fax: 619-453-2839

San Francisco
Phone: 415-955-2607
Fax: 415-397-6309

Santa Clara
Phone: 408-764-5000
Fax: 408-764-6740

Colorado
Englewood
Phone: 303-694-0674
Fax: 303-694-1670

Florida
Boca Raton
Phone: 407-392-0668
Fax: 407-362-4375

Maitland (Orlando area)
Phone: 407-661-1999
Fax: 407-660-0012

Tampa
Phone: 813-281-4635
Fax: 813-281-4619

Georgia
Atlanta
Phone: 404-395-2370
Fax: 404-395-2378

Illinois
Itasca
Phone: 708-250-5400
Fax: 708-250-5407

Massachusetts
Waltham
Phone: 617-466-9700
Fax: 617-890-9621

Michigan
Southfield
Phone: 313-262-1580
Fax: 313-262-1588

Minnesota
Minneapolis
Phone: 612-921-8335
Fax: 612-921-8327

Missouri
St. Louis
Phone: 314-984-6800
Fax: 314-984-6803

New York
New York City
Phone: 212-643-1733
Fax: 212-643-1738

North Carolina
Raleigh
Phone: 919-846-4875
Fax: 919-846-4877

Ohio
Cincinnati
Phone: 513-563-3559
Fax: 513-563-3523

Pennsylvania
Pittsburgh
Phone: 412-788-7740
Fax: 412-788-7744

West Conshohocken
Phone: 215-941-2777
Fax: 215-825-6296

Tennessee
Brentwood
Phone: 615-377-0754
Fax: 615-661-5587

Texas
Dallas
Phone: 214-980-5005
Fax: 214-980-5020

Houston
Phone: 713-864-3399
Fax: 713-864-4098

Utah
Salt Lake City
Phone: 801-264-6690
Fax: 801-264-6691

Virginia
Richmond
Phone: 804-273-0488
Fax: 804-273-0791

Vienna
Phone: 703-749-4200
Fax: 703-749-7988

Washington
Bellevue
Phone: 206-455-8530
Fax: 206-454-7402

Vancouver
Phone: 206-260-3483
Fax: 206-750-5542

3Com CANADA
Ontario
Ottawa
Phone: 613-782-2901
Fax: 613-782-2228

Richmond Hill
Phone: 416-882-9964
Fax: 416-882-9967

3Com Sales Offices Worldwide

To reach a 3Com office outside the country you are calling from, use the country code shown in parentheses and precede the number with the international access code of the country you are calling from (for example, 010 from the United Kingdom).

3COM HEADQUARTERS
P.O. Box 58145
5400 Bayfront Plaza
Santa Clara, California 95052-8145
Phone: (1) 408-764-5000
Fax: (1) 408-764-5001

3COM LIMITED
Hemel Hempstead, England
Phone: (44) 442 231000
Fax: (44) 442 236824

EUROPEAN OPERATIONS
France, Israel, Kuwait, North Africa, Saudi Arabia
3Com South Europe
Paris, France
Phone: (33) 1 69 86 68 00
Fax: (33) 1 69 07 11 54

Austria, Germany, Hungary Switzerland
3Com GmbH
Munich, Germany
Phone: (49) 89 62732 0
Fax: (49) 89 62732 233

Greece, Italy, Malta, Portugal, Spain
3Com Mediterraneo SRL
Milan, Italy
Phone: (39) 2 273 02041
Fax: (39) 2 273 04244

Belgium, Luxembourg, The Netherlands
3Com Benelux, B.V.
Nieuwegein, The Netherlands
Phone: (31) 3 402 55033
Fax: (31) 3 402 54630

Denmark, Finland, Iceland, Norway, Sweden
3Com Nordic AB
Stockholm, Sweden
Phone: (46) 8 632 9100
Fax: (46) 8 632 0905

England, Ireland, Scotland, Wales
3Com (U.K.) Limited
Buckinghamshire, U.K.
Phone: (44) 628 897000
Fax: (44) 628 897003

INTERCONTINENTAL OPERATIONS
Australia, New Zealand
3Com ANZA
North Sydney, Australia
Phone: (61) 2 959 3020
Fax: (61) 2 956 6247

Hong Kong
3Com Asia Limited
Hong Kong
Phone: (852) 868 9111
Fax: (852) 537 1149

Singapore
3Com Asia Limited, Singapore
Phone: (65) 538 9368
Fax: (65) 538 9369

Taiwan, Republic of China
3Com Asia Limited, Taipei
Phone: (886) 2 577 4352
Fax: (886) 2 577 4157

Japan
3Com Japan, Tokyo
Phone: (81) 3 3243 9234
Fax: (81) 3 3243 9235

3Com Middle East
P.O. Box 52898
3rd Floor, Al Moosa Tower
Dubai-Abu Dhabi Highway
Dubai, United Arab Emirates

Latin America
3Com Corporation
Santa Clara, California
Phone: (1) 408-764-6462
Fax: (1) 408-764-5742

México
3Com de México
Ejército Nacional No. 539-12
Col. Granada, CP 11520, México DF
México
Phone: (011) 525-531-0591
Fax: (011) 525-254-3159

Scorecard!

3TECH Scorecard

Please take just one minute to "grade" this issue and fax back to *3TECH!* Which of the articles and departments do you rate the highest and which the lowest?

Please rate the following articles and departments, 1 = needs improvement, 5 = excellent, give us more of the same!

Rating	Feature Articles
______	High-Performance Scalable Networking with Routed ATM—Enhancing the Performance of Today's LAN Technologies
______	Transcend Network Management Solution—Integrated Management Based on Connectivity Systems and SmartAgent Software
______	3Com Leads the Multivendor APPN Pack—Implementing IBM's Advanced Peer-to-Peer Networking Code on the NETBuilder II Platform
	Departments
______	Tech Tips
______	Q&A: What Is TUBA?
______	Bug Reports
______	3Info
______	New Products at a Glance

Inside and outside U.S./Canada
Mail this form to:
3Com Corporation
Attn: Linda Webb
P.O. Box 58145
Santa Clara, CA 95052-8145

Or fax this form to:
408-764-6477
Attn: Linda Webb

Calendar | 3TECH Subscription Information

Upcoming Events

January 16–20
16th Annual Pacific Telecommunications Conference
Honolulu, Hawaii

February 2
ICC Link
Milan, Italy

February 9
ICC Link
Madrid, Spain

February 15–17
Networks Expo
Boston, Massachusetts

March 16–18
IS Quest '94
Santa Clara, California

March 16–23
CeBit
Hanover, Germany

For information on 3Com seminars in your area, call your local 3Com sales office.

3TECH Subscription Information

☐ **Change of Address**
I already receive *3TECH,* but please make the following corrections.

☐ **Subscription Cancellation**
Please cancel my *3TECH* subscription.

PLEASE PRINT:

Name ______________________

Title ______________________

Company ______________________

Mailing address (street number or P. O. box number) ______________________

City/state or province | Zip+extension or postal code | Country

Country/area code | Telephone number | Country/area code | Fax number

Inside and Outside U.S./Canada
Mail this form to:
3Com Corporation
Attn: Linda Webb
P.O. Box 58145
Santa Clara, CA 95052-8145

Or fax this form to:
408-764-6477
Attn: Linda Webb

Bulk Rate
U.S. Postage
PAID
Sunnyvale, CA
Permit No. 5

ADDRESS CORRECTION REQUESTED

3Com Corporation
P.O. Box 869038, Plano, Texas 75086

Printed in USA 740000-001