Local Ethernet Bridges

A Report from NSTL

In this report:

Synopsis

Focus
This report describes how bridges improve network functioning, explains how to choose a bridge product, and compares the relative merits of six local Ethernet bridges.

Products Tested

HSB-EE
CrossComm Corp.

ProBridge 8033
Hughes LAN Systems

EtherMaster 100
Netronix, Inc.

Intersect
Persoft, Inc.

INX400/L
Racal InterLan

4660
Retix

Product Recommendations

- Retix 4660
- CrossComm HSB-EE
- Netronix EtherMaster 100

Source
Based on data generated by tests designed and conducted by National Software Testing Laboratories, Inc. (NSTL), a division of Datapro Research Group, Plymouth Meeting, PA 19462. Telephone (800) 223-7093.

Ratings Key (On a scale of 0 to 10)
Ratings
● 7.0 - 10.0
○ 5.0 - 6.9
⊗ under 5.0
★ Recommended

Overall Rating	Product Name	Performance	Features	Usability	Price
9.5	Retix 4660	●	●	●	$3,750 ★
8.5	CrossComm HSB-EE	●	●	●	$4,350 ★
7.8	Netronix EtherMaster 100	●	●	●	$3,400 ★
7.3	Racal InterLan INX400/L	○	●	●	$3,195
7.0	Hughes ProBridge 8033	○	●	●	$2,995
5.0	Persoft Intersect	○	⊗	●	$1,495

Overview

As any organization grows and changes, its local area network undergoes similar growth. At its simplest level, a network grows through added resources (workstations, servers, printers), but extended growth more often requires connecting multiple networks. Internetworking addresses the problem of connecting potentially dissimilar hardware, software, and network platforms into a consolidated "enterprise-wide" network.

The growing need for internetworking capabilities has spurred the development of a multitude of products. According to a study by Market Intelligence Research Corp. (MIRC), sales of internetworking products such as bridges, routers, and gateways increased at a compound annual growth rate of 83 percent from 1987 to 1989, and 1991 sales projections for various internetworking products range from $186 million to $196 million. Bridge vendors indicate that while the market for token-ring bridge products is growing rapidly, Ethernet bridge sales overall are greater, benefiting from a larger installed base.

Evaluation Criteria

As part of NSTL's commitment to provide businesses with the tools and information needed to understand and evaluate internetworking products, this evaluation focuses on IEEE 802.3-compatible bridges capable of bridging two Ethernet segments at the Media Access Control (MAC) layer of the ISO model. MAC-layer bridges filter or forward packets based on MAC-layer information in the packet; specifically, the source address, destination address, and frame size information contained within the first 14 bytes.

NSTL examined the bridges' installation, configuration, adherence to ISO standards, performance, and support for an extensive list of features. The evaluation emphasizes the products' performance when given the task of bridging two network segments running Novell NetWare 386 Version 3.1.

RAD Network Devices declined participation in the evaluation, citing its greater emphasis on remote bridges. Ungermann-Bass, 3Com, and ACC evaluation units were not available within the testing time period.

Filtering/Forwarding

Vendors of local Ethernet bridges describe bridge performance in terms of Ethernet frame (or packet) filtering and forwarding rates. Specifications listing rates of 13,000 or 14,000 packets per second (pps) imply that the bridge can handle just about any level of network traffic on an IEEE 802.3-compatible network. Unfortunately, latency (the time elapsed while the bridge receives and retransmits a packet) slows these rates. Assuming a bridge has room to store incoming packets while other packets queue for transmission, only a small initial delay occurs as the first small group of packets is readied for transmission. After this initial delay, packets stream at the specified rate. Latency delays are typically less than a few milliseconds.

On some networks (Novell NetWare 386 in particular), all data packets sent by an application between a server and workstation require acknowledgment frames. Cumulative packet latencies at the application level introduce real performance implications. A bridge slows applications such as Lotus 1-2-3 running on a Novell network when a bridge is the intermediary between a workstation and server (confirmed by NSTL's performance test configuration).

Bridge Applications

Traffic Isolation

As a network grows, more and more information must pass through the physical media connecting network resources. At some point, the level of traffic on the network will degrade overall network performance, making further expansion unrealistic. Existing networks may experience this type of loading during periods of heavy use.

A local Ethernet bridge dividing one large Ethernet network segment into two connected Ethernet segments can effectively isolate local traffic (source and destination addresses on the same

Figure 1.
Bridged Network Topologies

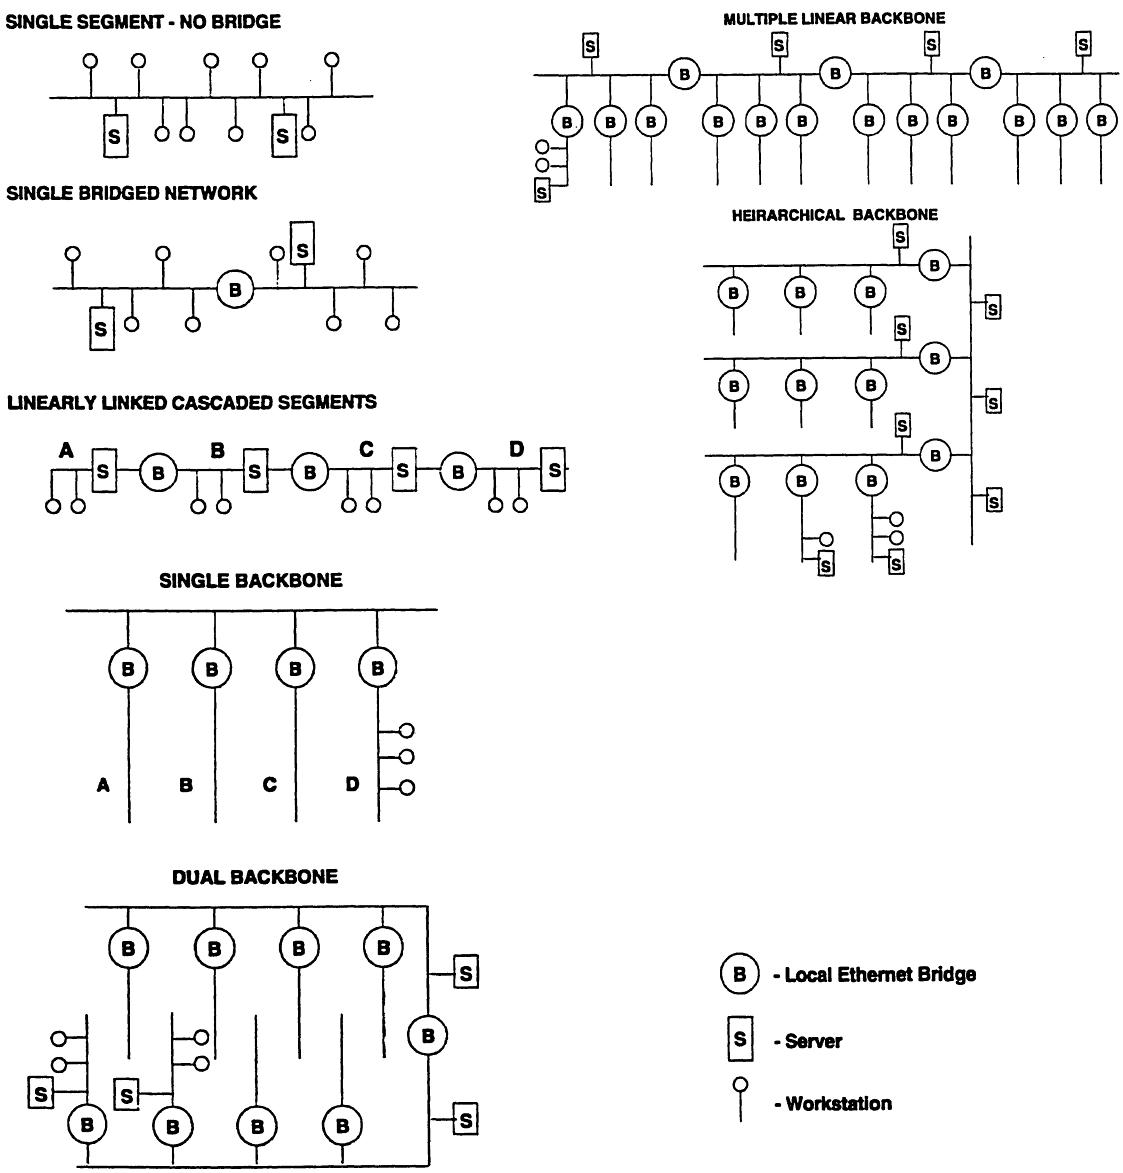

segment). Without a bridge, traffic originating anywhere on the network appears at all other addresses, effectively using up a portion of the available network bandwidth. A bridge eliminates nonlocal traffic and can potentially enhance performance on each segment.

Optimize Bandwidth Utilization

When more than three or four Ethernet network segments are connected linearly by bridges, the volume of intervening traffic on the middle segments may interfere with communications between the outermost segments. Figure 1 shows bridged network topologies that address traffic overloading. Using a single segment as a backbone for additional bridged segments eliminates local traffic from the backbone. Network organization and management determine what percentage of each segment's traffic will be nonlocal. As more segments are added, traffic may overload the backbone. Further expansion may require dual backbones connected linearly with a bridge. Each backbone is effectively isolated from the local traffic of the other backbone's segments. A hierarchical backbone scheme supports even more expansion with a backbone for backbones. A hierarchical backbone avoids the pitfall of multiple linear backbones (i.e., intervening traffic on middle backbones can become too heavy).

The end result of these and other topology changes is to effectively expand the total available bandwidth for the extended network. Beyond a certain level of complexity, clever topologies alone may not provide adequate solutions to traffic overloading. Using high-speed media such as fiber optic cable for network backbones enables upgrades to higher bandwidths.

Fault Tolerance

Bridges help network administrators isolate network faults. In addition to providing electrical isolation between segments, multiple bridges can provide routing fault tolerance with redundant network connections. If an active bridge fails, a backup or standby bridge can be activated to ensure communications between network segments.

Load Balancing

A well-planned traffic isolation scheme uses one or more bridges to balance the load of network traffic on interconnected network segments. Priorities can be assigned to bridge channels and to individual bridges.

Enhanced Security

Bridges can function as extended network connection control mechanisms, adding to existing network security. General security measures such as protocol filtering are easily implemented for some bridges. Protocol filtering controls connections based on packet protocol type. Most bridges can control frame filtering and forwarding to specific network addresses. With proper administration, a bridge can be set up to deny one or more nodes access to a network segment. In a large extended network, such administration may not be easy. All the bridges "learn" the network addresses of connected devices, and a network with several hundred or several thousand learned nodes might require extensive manual editing of bridge address tables.

Mixed Media Connection

Most local bridges can connect network segments with different physical media. All the test bridges can connect thin and thick Ethernet segments. The Persoft, Racal InterLan, and Retix also connect to twisted-pair Ethernet. The Retix provides a modular interface port for connecting to thin Ethernet, thick Ethernet, unshielded twisted-pair Ethernet, and fiber optic Ethernet media.

Physical Media Extension

Like repeaters, bridges receive network frames and retransmit them, providing a degree of electrical conditioning of the transmitted frame's signal. Electrical conditioning combined with topology options can greatly extend the media lengths allowed with a single network segment.

IEEE 802.3 Ethernet Standard

The Institute of Electrical and Electronics Engineers (IEEE) 802.3 Ethernet Standard (Version 2) specifies a minimum 60-byte Ethernet frame, excluding an 8-byte preamble and 4-byte frame check sequence (FCS). The minimum frame length ensures that a frame's duration is longer than the worst-case propagation delay of 45 milliseconds through the network. A frame includes a 6-byte destination address, a 6-byte source address, and a data field of at least 48 bytes. The first two data field bytes describe the length of the data.

To determine whether to filter or forward a packet, the test bridges examine the first 14 bytes (of the 60 bytes): the destination address, source address, and 802.3 data length. These 14 bytes (collectively known as the Data Link Control Header in the Network General Sniffer terminology) are the basis for bridge functionality at the Data Link level of the International Organization for Standardization (ISO) model. (NSTL used a Network General Sniffer to verify Ethernet traffic levels and to assist in performance testing.)

ISO OSI Model

Bridges belong to a family of internetworking devices that includes repeaters, bridges, routers, and gateways. The ISO Open Systems Interconnection

Figure 2.
ISO OSI Layers

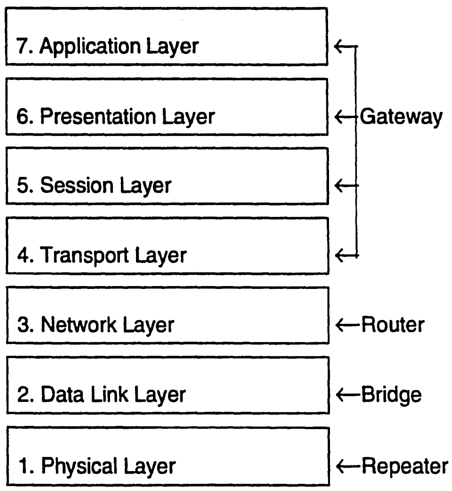

(OSI) Reference Model describes the communications layer at which these devices operate (see Figure 2).

The distinctions among repeaters, bridges, routers, and gateways are not always clear because some products do not operate in an ISO environment. Furthermore, manufacturers add routing functions to bridge products (routing bridges) and bridging functions to router products (brouters). Higher performance levels are gained with internetworking devices that do not have to examine higher layer information within a packet to determine a course of action. Consequently, the most efficient devices are those that depend only on the lowest layer of the ISO model. The distinction between routers and bridges is less important than the distinction between the layer of the ISO model at which a device operates.

Repeaters (Physical layer) provide conditioning, retransmission, and electrical isolation of internetwork signals. Essentially, a repeater forwards or repeats electrical data transmissions from a source network segment to a destination network segment.

Bridges operate at the MAC (Media Access Control) sublayer of the Data Link layer. Bridges permit network traffic isolation, connect separate network segments into a single logical network, regulate traffic by examining the destination address of each packet, and are transparent to protocols higher than the Data Link layer, such as Novell's IPX/SPX, Banyan VINES IP/SPP, XNS, TCP/IP, IBM SNA, and X.25 protocols.

Routers (Network layer) connect network segments running the same protocol. Packets can be filtered, forwarded, or routed based on address and protocol information contained within the packet. Routers provide multiple traffic paths to a given destination with intelligent selection of optimal (e.g., lowest cost, fastest, most reliable, etc.) paths through an internetwork. Routers that provide simple bridging capabilities (brouters) can route one or more protocols and bridge all other traffic.

Gateways combine the functions of repeaters, bridges, and routers and can translate data packet messages between network segments with different protocols. Gateways operate at the Transport level of the ISO model and higher, depending on the level of compatibility of the connected protocols.

Learning Bridges

Conventional bridges (also called learning bridges) learn and keep track of the source addresses of transmitted packets. Based on comparisons with learned addresses, the bridge filters or forwards packets by examining their destination addresses. Filtered traffic remains local rather than being forwarded to a nonlocal segment. Packets can be passed regardless of protocol information because the bridge examines only the MAC-layer address information contained in the packet.

When several bridges connect several Ethernet segments, traffic loops can develop between segments in which frames are retransmitted indefinitely by the various bridges on the loop. Learning bridges do not automatically avoid traffic loops in the Ethernet topology, so the network designer/administrator must correctly segment an extended network. With multiple active bridges between network segments, the administrator can balance network loading by controlling priorities and address table entries for each bridge.

Routing Bridges

Routing bridges offer intelligent path selection and controlled access to network resources that can be helpful in managing large, complex networks. Routing bridges feature a spanning-tree algorithm, custom filtering, and source-explicit forwarding.

Spanning-Tree Algorithm

Based on an IEEE draft standard, the spanning-tree algorithm permits multiple paths (and consequently, loops) between any two network segments on a bridged network. Using administrator-defined path costs, the spanning-tree algorithm automatically selects an optimal path between segments. If a bridge fails and an active path becomes disabled, the algorithm reconfigures the network to use alternative paths. The spanning tree automatically configures large, complex network topologies and easily accommodates topology changes by relearning addresses and activating bridges in standby mode. Redundant links between segments (through the use of backup bridges) can enhance network reliability and fault tolerance.

The spanning-tree algorithm for Ethernet bridges prevents traffic loops by assigning bridge priorities, setting some bridges to standby mode, and defining one bridge as the "root" of the spanning "tree." The "tree" hierarchy can lead to very heavy traffic through the root bridge, and the "no loop" requirement prevents some types of load balancing. Spanning-tree bridges communicate with other spanning-tree bridges using IEEE 802.1 Bridge Protocol Data Units (BPDUs) called "hello" messages.

Custom Filtering/Forwarding

With custom filtering, a bridge filters and forwards packets based on a set of conditions. The administrator can designate that packets of specific length or protocol type, or all broadcast packets, be always filtered or always forwarded. Custom filters enhance a bridge's traffic isolation capabilities.

Source-Explicit Forwarding

Source-explicit forwarding designates that specific network addresses forward packets through a specific bridge port and prevents other addresses from forwarding through that port. By limiting access in this way, administrators can gain more control over the security of specific network segments or resources.

Source-Routing Bridges

Source-routing bridges (which were not evaluated) determine the appropriate network path based on an additional routing field contained in the packet's MAC layer. Source control routing bridges are almost exclusive to IEEE 802.5 token-passing rings.

Source-Routing Transparency

An emerging IEEE standard, source-routing transparency (SRT) specifies rules for bridge design and functionality that require interoperability of spanning-tree (802.1 Network layer) and source-routing bridges (IBM's 802.5 MAC layer). With the advent of bridges that connect multiple networks with different topologies (e.g., Ethernet and token-ring), this type of standard may become more important in bridge design.

Performance

NSTL tested bridge performance on a Novell NetWare 386 Version 3.1 network. Applications were installed on the network from a Maynard MainStream 1300 DAT tape backup unit. Performance tests measure application performance with and without background traffic. Other tests measure the delays associated with transmitting frames of various sizes across the bridge.

Without background traffic, the application tests create relatively light network loading conditions with relatively few frame collisions and fragments. Under these conditions, bridge performance is mostly a function of the delay (latency) associated with forwarding each frame across the bridge. The per-frame delay under light network loading is relatively independent of the exact load percentage. Because NetWare 386 requires an acknowledgment for every frame, the delays are cumulative.

The Retix, Netronix, and Persoft exhibit relatively linear relationships between latency and packet size. Bridge latencies are measured for frame sizes from 100 to 500 bytes, and latencies for larger frame sizes can be extrapolated using the linear relationships. Cumulative delays for the Lotus 1-2-3 test can be calculated using frame size distributions for the Lotus test (measured with the Network General Sniffer) and measured frame latencies. Adding the calculated delays to the No Bridge times for this test accurately predicts the completion time with each bridge. The high level of predictability validates the assumption that latency is primarily a function of packet size rather than traffic intensity.

Summary

The Retix, CrossComm, and Netronix bridges successfully minimize cumulative packet delays with a broad range of packet sizes and traffic conditions. All the bridges support the spanning-tree algorithm, learn source addresses, and automatically build source address tables. The products support address tables of varying sizes and differ in their capability to save source addresses in nonvolatile RAM.

The Retix excels in network scenarios where heavy traffic is expected and packet delays must be minimized. The Retix bridge can be ordered from the manufacturer configured with either SNMP or MAP/TOP management capabilities. With SNMP management capabilities, the bridge can be managed entirely through its console port. The optional 211 Local Bridge Manager ($1,450) requires that a MAP/TOP management image be loaded. The MAP/TOP image can be loaded using the 4660/4660-S Conversion Kit ($250). The high-end Microsoft Windows-based 5010 Network Manager Center sells for $5,500.

The CrossComm supports a very large source address database and performs very well with heavy traffic. Its Internetworking Management Software (supplied) can save and retrieve modified source address table entries and supports IBM's NetView.

The Netronix offers good price/performance with heavy traffic and provides very good source address table support. It is SNMP compatible, and its EtherMaster Support Software supports bridge management from any network node.

The Racal InterLan provides good performance under light to moderate traffic conditions, SNMP support, and very good configuration and administration facilities. It comes with excellent documentation and is easy to install. The Racal InterLan performs better with light background traffic than with no traffic; in fact, it required light background traffic in order to complete NSTL's latency tests.

The Hughes comes with excellent documentation, can operate as a simple learning bridge, provides robust configuration and administration features, and provides a good price/performance ratio for networks with light to moderate traffic. At $2,995, the Hughes is among the least expensive of the test bridges.

The Persoft offers a relatively inexpensive alternative to high-end bridges for networks with moderate traffic and throughput. Performance is subject to the limitations of the system processor and bus architecture.

Product Evaluations

CrossComm HSB-EE

Product Summary

- Very good performance
- Latencies increase at 400-byte packet size without traffic
- Short latencies under heavy traffic conditions
- Comprehensive documentation
- Automatic learning of source addresses
- Up to 40,000 source address table entries
- Fast proprietary address lookup engine
- Can save 128 addresses in nonvolatile RAM
- Easy installation
- Good diagnostics
- Internetworking Management Software (IMS)
- Spanning-tree algorithm
- SNMP support

Hughes ProBridge 8033

Product Summary

- Excellent documentation
- Automatic learning of source addresses
- Up to 3,000 source address table entries
- Can save 800 addresses in nonvolatile RAM
- Spanning-tree algorithm in Version 1.1 software
- Can operate as simple learning bridge
- Very good source address table
- Robust configuration/administration features
- Good performance under light to moderate loads
- Good price/performance excluding extremely heavy traffic
- External power supply

Netronix EtherMaster 100

Product Summary

- Very good performance
- Small latencies at all packet sizes
- Easy installation
- Good diagnostics and front-panel indicators
- Automatic learning of source addresses
- Up to 8,000 source address table entries
- Can save 1,000 addresses in nonvolatile RAM
- Very good source address table
- EtherMaster support software
- Spanning-tree algorithm
- Very good price/performance excluding heavy traffic
- SNMP support

Persoft Intersect

Product Summary

- Inexpensive alternative excluding heavy traffic and high throughput requirements
- Easy software installation
- Two Ethernet adapters install in host
- Potential performance increase in more powerful hosts
- Adequate performance with light to moderate traffic
- Subject to ISA bus limitations
- Automatic learning of source addresses
- Monitors address database to prevent filtering/forwarding
- Cannot store addresses in nonvolatile RAM
- Relatively long latencies for all packet sizes
- Fails all benchmarks with intense IPX traffic

Racal InterLan INX400/L

Product Summary

- Excellent documentation
- Easy installation
- Automatic learning of source addresses
- Up to 8,000 source address table entries
- Can save 50 addresses in nonvolatile RAM
- Good performance with light to moderate traffic
- Faster with light background traffic than without traffic
- Consistently small latencies; requires light background traffic
- Fails latency test with intense IPX traffic
- Spanning-tree algorithm
- Very good configuration/administration features
- SNMP support

Retix 4660

Product Summary

- Excellent performance
- Very short forwarding latencies
- Automatic learning of source addresses
- Up to 8,192 source address table entries
- Can store 4,000 addresses in nonvolatile RAM
- Good management software

- Spanning-tree algorithm
- SNMP support
- Robust configuration/administration facilities
- Supports fiber optic media
- Very good high-end price/performance

Product Recommendations

Retix 4660

Excellent performance earns the Retix 4660 NSTL's recommendation. The Retix excels in network scenarios where heavy traffic is expected and packet delays must be minimized. Used with network operating system protocol schemes that generate send/acknowledgment packet pairs for every transmission, the Retix will minimize the cumulative effect of packet delays for a broad range of packet sizes and traffic intensities.

The Retix' modular Network Interface Cards (NICs) slide easily into place and enable on-site configuration for thin, thick, twisted-pair, and fiber optic Ethernet connections. Compatibility with fiber optic media offers an upgrade path to 100M bps FDDI. Administration functions are accessible through menu-driven bridge management software, and network management protocol support includes SNMP and CMIP. Spanning-tree protocol with automatic loop detection, a large address database, and the capability to save a large number of addresses in nonvolatile RAM add versatility in bridging local Ethernet segments.

CrossComm HSB-EE

For installations with large numbers of workstations, the CrossComm provides the largest source address database and good performance with heavy traffic. Easy installation, comprehensive documentation, informative front-panel indicators, and good source address table support add to the product's attractiveness. An extremely large address table and fast look up engine make the CrossComm a good choice for extended networks with large numbers of nodes. Under any traffic conditions, the CrossComm performs well and minimizes cumulative packet delays for network protocols requiring send/acknowledgments for every transmission. It supports the spanning-tree protocol, SNMP, and NetView and automatically learns network addresses. Manually entered and modified address table entries can be saved and retrieved from DOS files.

Netronix EtherMaster 100

Good performance and very consistent operation, excellent source address table support, robust configuration and administration features, and easy installation recommend the Netronix as a lower-priced alternative to the CrossComm and Retix. Spanning-tree protocol, a large address database, the capability to save many addresses in nonvolatile RAM, and the capability to minimize cumulative packet delays add to the Netronix' overall value.

Rating Summaries

Figure 3.
CrossComm HSB-EE Ratings

Figure 4.
Hughes ProBridge 8033 Ratings

Figure 5.
Netronix EtherMaster 100 Ratings

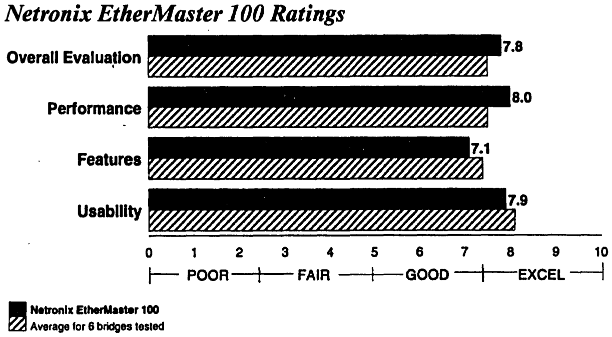

Figure 6.
Persoft Intersect Ratings

Figure 7.
Racal InterLan INX400/L Ratings

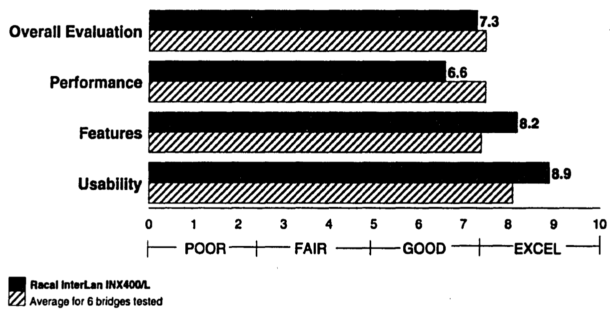

Figure 8.
Retix 4660 Ratings

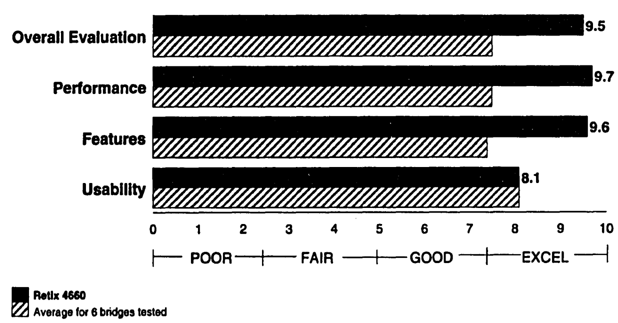

Overall Evaluation

Heavy traffic situations and potentially complex network management scenarios require bridges with strong performance and bridge management tools. The Retix easily outscores the other bridges with the best performance and features and very good usability. The methodology's heavy reliance on performance accentuates differences between the Retix and the other bridges. The Retix performs flawlessly, even with extremely heavy background traffic. Its support for FDDI-compatible media enables its use with 10M bps fiber optic Ethernet media.

The CrossComm and Netronix offer similar levels of performance, features, and usability. Netronix provides excellent source address database management. CrossComm offers better overall configuration and administration functions. The CrossComm and Netronix perform extremely well with heavy background traffic.

The Racal, Hughes, and Persoft experience some problems with heavy network traffic. The Racal performs very well under light and moderate traffic conditions and provides the most features and excellent usability. The Hughes shows much slower performance under heavy traffic. Its flexible configuration options meet almost any need; it comes with good documentation and installs easily.

Relatively few configuration options and features make the Persoft very easy to use. In the test system, its performance is adequate under light network loading and poor under heavy traffic. In a more powerful system, the Persoft should perform faster, but gains may be offset by limitations of the ISA bus architecture. The product lacks any source address table management software.

Methodology

The Overall Evaluation is a weighted average of scores for the individual criteria.

Overall Evaluation Score = (6 x Performance Score) + (3 x Features Score) + (Usability Score) ÷ 10.

Performance

NSTL's benchmark suite measures performance differences with varying levels of network traffic. Heavy traffic accentuates relative performance differences and causes some products to falter and fail. Internal bridge latency, the inherently short transmission delay when forwarding a data packet, varies with packet size and can be as important as packet throughput to overall bridge performance. Latency effects for the local Ethernet bridges operating in a Novell NetWare 386 environment are cumulative and directly affect the application benchmark performance. A bridge's processor speed, memory, address database size and functions, packets queued and spacing, and algorithm efficiency also affect performance.

The Retix uses a 20MHz 68020 processor and 200K-byte frame buffer. Its internal forwarding latency is the shortest at almost all packet sizes, and it does not drop packets under NSTL's testing. Small delays under all traffic conditions confirm the importance of processor type and speed in good performance.

The CrossComm uses a 16MHz Intel 80376 processor and 128K-byte ultrahigh-speed RAM frame buffer, which accounts for its smaller packet delays under heavy traffic conditions. Without traffic, the CrossComm's latency increases substantially with larger packet sizes; with IPX traffic, latencies remain very short for all packet sizes.

The Netronix uses a 16MHz Intel 80376 processor and 256K-byte frame buffer, but its internal latencies are longer than the Retix or CrossComm. The Hughes uses a slower 16MHz 386SX processor and produces internal latencies similar to the Netronix'. The Hughes performs poorly under heavy IPX traffic and drops packets in NSTL's test. The Racal outperforms the Hughes despite its less powerful 10MHz 80286 processor, in part because of its faster performance under heavy IPX traffic.

The Persoft installs in a host PC and is subject to performance limitations imposed by the host. Transfer of Ethernet packets across the PC bus, subsequent address table lookups, filtering/

forwarding, and possible retransmission of forwarded packets contribute to packet delays. Given the Persoft's 8-bit Ethernet adapters, forwarding a 500-byte packet across a typical 8MHz ISA bus takes a minimum 0.4 to 1.0 millisecond, depending on the number of additional wait states. (The Retix' delay with 500-byte packets is 0.46 millisecond.) A system with a more powerful processor will not eliminate the limitation of the 8-bit ISA bus interface. During testing, the Persoft dropped more than three times the number of packets dropped by the Hughes (only these two products dropped packets).

Methodology

For each aspect of the benchmarks, individual performance scores are weighted and scaled. The overall performance score represents a weighted average of performance indexes for all benchmarks.

Performance Score = (10 x Latency Score) + (4 x Lotus 1-2-3 Score) + (4 x Microsoft C 5.0 Score) + (4 x Foxbase+/LAN Score) + (4 x Xcopy to Server Score) + (4 x Xcopy from Server Score) + (Dropped Packets Score) ÷ 31.

Features

The Retix' high features rating results from its good support for different media, physical connections, and IEEE 802.3 standards; good diagnostics and front-panel indicators; and good support for network management protocols. The Retix, Hughes, and Racal provide thorough and flexible configuration and administrative features; the Hughes earns high marks for protocol-specific filtering and its source address database.

The Hughes currently comes with Version 1.0 of its bridge management software. Version 1.1 (now shipping) adds a spanning-tree algorithm and Permit Only protocol filtering, which enables the bridge to selectively filter protocols. With Version 1.1, spanning-tree bridges can coexist on an extended network with non-spanning-tree bridges. Configuring the Hughes as a spanning-tree or simple learning bridge enables better load balancing by avoiding root bridge traffic congestion.

Methodology

The features rating is a weighted average of scores for the individual criteria. A complete list of features and their methodology weights appears in Table 4.

Usability

Bridge usability relates directly to the availability and quality of supplied utilities and configuration/administration options.

Racal provides well-organized Installation, Manager's, and Command Reference manuals. The Racal setup requires installation of a software cartridge, and access to the bridge console requires a serial cable connecting the console port and a terminal or system running a terminal emulation program.

The Hughes manual provides detailed, comprehensive descriptions of configuration and administrative commands and a good index and table of contents. Software installs from a cartridge. The Hughes bridge lacks a power switch, and its external power supply is awkward.

The Retix, Netronix, and CrossComm install easily. The Retix requires a serial connection to access its console; the Netronix and CrossComm console functions can be accessed across the LAN. The manuals are fairly comprehensive but lack overall organization and indexes.

The Persoft bridge requires installation of two 8-bit Western Digital Ethernet cards in a dedicated host PC. The Ethercard Plus Installation Guide supplied with the adapters furnishes sufficient information to accommodate changes to the I/O base address, IRQ request, and network selection settings. The setting changes complicate installation, but the Persoft's few console features make it simple to operate.

Methodology

The usability rating is a weighted average of scores for the individual criteria.

Usability Score = (3 x Physical Installation Score) + (2 x Console/Administration Hookup Score) + (3 x Console/Administration Use Score) + (2 x Manual Organization Score) + (2 x Manual Clarity Score) + (3 x Manual Comprehensiveness Score) ÷ 15.

Performance Results

Local Ethernet bridge performance, characterized by maximum packet filtering and forwarding rates, relies on measurements made by dedicated test equipment in a controlled environment. Such "raw" measurements do not address local Ethernet bridge behavior in an active extended network context with unique traffic situations. Packet size, queuing methodology, interpacket spacing, bandwidth utilization, and internal bridge latency play important roles in the actual performance of an Ethernet bridge connecting two or more local networks.

Test Configuration

To test local Ethernet bridges, NSTL set up a Compaq Deskpro 386/20 server running Novell NetWare 386 Version 3.1 with a Compaq Deskpro 386/25 primary workstation, five Everex 386/SX secondary workstations, and a Compaq 386/20 controller workstation that automates the tests. A Micro Express 486/33 acted as a receiver for node-to-node IPX conversations with the primary workstation, and another Compaq Deskpro 386/25 generated intense IPX traffic on the network. At all times, a Network General Sniffer was active on one network segment. All test data, application programs, and custom program files residing on the primary workstation were archived using a Maynard MainStream 1300 DAT Tape Backup System.

The test network was configured with two segments connected by a local Ethernet bridge. One segment (Subnet 0) consisted of the server, IPX Receiver, and Sniffer connected via thin Ethernet cable. The other segment (Subnet 1) included the five secondary workstations, the primary workstation, and IPX traffic generator connected with thin Ethernet cable. The server, IPX receiver, primary workstation, controller workstation, and the IPX traffic generator used 16-bit Western Digital Ethercard Plus 16 adapters, and the five secondary workstations used 8-bit Ethercard adapters. Both segments were terminated with terminating resistors; Subnet 1 used a grounded terminating resistor. The Sniffer connected to the network using a proprietary 16-bit Ethernet card supplied by Network General and ran Version 3.0 of Network General's Ethernet Analyzer and Ethernet Monitor Software.

Persoft Intersect

Persoft's bridge product includes two 8-bit Ethernet cards and bridge software. Its performance is subject to the limitations of the host architecture and performance characteristics. NSTL tested the Persoft in an Everex 386/SX as representative of relatively powerful and low-cost workstations typically dedicated to bridging functions.

Application Tests

Four benchmarks measure local bridge performance with Lotus 1-2-3 Release 3, Microsoft C, FoxPro, and Xcopy operations. Each test runs from the primary workstation with the secondary stations idling (no traffic) as a baseline measurement. Background traffic scenarios contrast with the baseline. Traffic scenarios include background application traffic generated at five secondary workstations by Foxbase+/LAN and Lotus 1-2-3 and intense IPX traffic generated by the Compaq Deskpro 386/25 (with no traffic at the workstations).

Internal Latency

Internal latency refers to the time the bridge takes to process an Ethernet packet and forward or filter the packet based on destination address information in the packet's Data Link Control Header. Upon receiving a packet, the bridge references its internal database of Ethernet node addresses to determine the location of the packet's destination address. A fairly small internal address database involves few address comparisons and minimal bridge processing. If the destination is located on the same segment as the source address, the bridge filters the packet, which remains local. Packet transmission and receipt on the local Ethernet segment proceed independently of bridge activity. If the packet's destination address is on a network segment other than its source, the packet is forwarded. Forwarding packets (analyze and retransmit) takes longer than filtering (analyze and do nothing), and packets queued in the bridge's memory for forwarding contribute to retransmission delay or latency.

NSTL's latency tests measure the time required to forward a large number of packets across the bridge. Two network nodes residing on different segments conduct an IPX conversation across the bridge using NSTL's custom send and receive programs, which are optimized for throughput with no disk access or extraneous data processing. The IPX send program generates address-specific packets, and the receive program checks for dropped packet time-outs to ensure successful forwarding of all packets. Each conversation (9,999 packets) is repeated using five packet sizes (100 to 500 bytes), and all conversations with all packet sizes are run without traffic and with intense IPX traffic. Latency results are compared with times for conversations generated across the network set up as a single segment.

Background Traffic

Lotus 1-2-3 and Foxbase+/LAN generate high-level background traffic on the secondary workstations. Foxbase performs a transaction no more than 30 times a minute; the transaction looks up indexed records in three files and locks, updates, and unlocks one record in each file. Lotus executes a macro that combines and saves files metered to be active no more than every 20 seconds on each station. As configured in the test bed and verified using the Network General Sniffer, Lotus traffic uses 1 to 5 percent of the network bandwidth; Foxbase traffic uses from 2 to 6 percent of the available bandwidth.

An NSTL IPX program generates low-level IPX background traffic optimized for throughput; the program continuously sends 72-byte Ethernet packets to a nonexistent network address to ensure automatic forwarding across the bridge and to generate intense traffic on both segments. IPX background traffic results in steady network utilization at about 36 percent of the network bandwidth.

Lotus 1-2-3

A Lotus macro is invoked at the primary workstation to load and save a 3M-byte spreadsheet, moving a substantial amount of data across the network. This test provides a well-balanced data transmission scenario between the workstation and server. The test runs without traffic, with Lotus 1-2-3 and Foxbase traffic, and with intense IPX traffic on both segments.

Figure 9.
Lotus 1-2-3

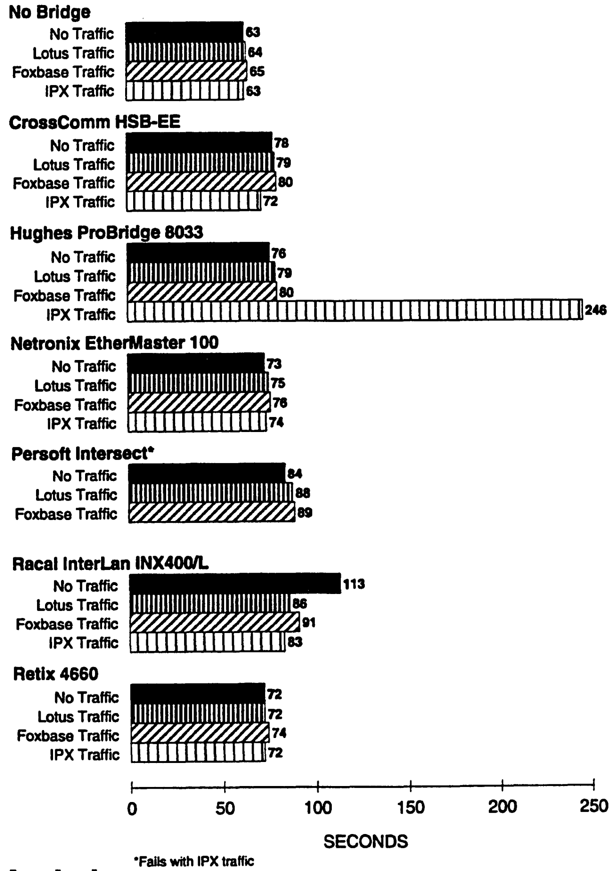

Analysis

The intense IPX traffic utilizing approximately 36 percent of total network bandwidth presents problems for the Persoft and Hughes bridges (36 percent bandwidth utilization is uncommon, but it highlights bridge performance under duress).

The Persoft bridge (in the test system configuration) fails completely with IPX traffic. The primary workstation cannot be initialized by the control station. The level of packet collisions, resultant exponential back-off, bridge buffering demands, and diminished available bandwidth can exceed the Hughes bridge's forwarding capabilities. Both the Hughes and Persoft perform well with other types of traffic.

The Racal bridge performs better with traffic than without. The Racal's internal latency results (see Internal Latency section) help explain this anomaly.

Microsoft C 5.0

Microsoft C compiles and links a large ensemble of 25 Xlisp 2.0 C source code files (10,928 lines;

Figure 10.
Microsoft C 5.0

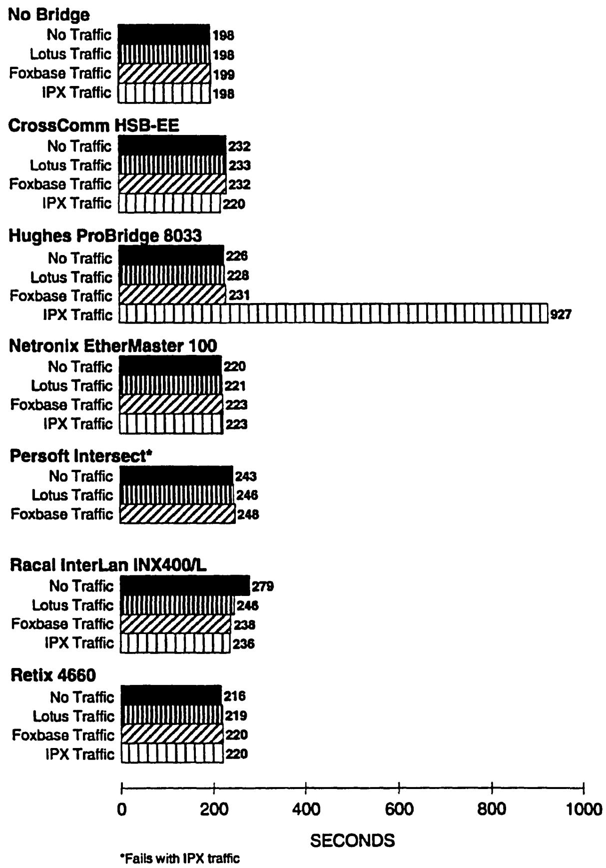

245,222 bytes), forming an executable file of 147,228 bytes. The total amount of data downloaded from the server is far larger than the amount of data traveling from workstation to server. The test runs without traffic, with Lotus 1-2-3 and Foxbase traffic, and with intense IPX traffic on both segments. This test primarily measures data transfer speed from the server to the primary workstation; source code is loaded and compiled at the primary workstation.

Analysis

The Retix performs consistently well in all versions of the Microsoft C tests (only about 9 to 11 percent slower than transmission with no bridge).

The Racal performs at its slowest with no traffic and at its fastest with intense IPX traffic. Without traffic, the Racal is about 40 percent slower than without a bridge and about 29 percent slower than the Retix; these margins narrow with traffic. With IPX traffic, the Racal is only around 18 percent slower than no bridge and the Retix. (IPX traffic in the internal latency test causes the bridge to fail; for an explanation, refer to the analysis for the Internal Latency section.)

IPX background traffic apparently overwhelms the Persoft's capability to process incoming packets. With no traffic and with Lotus and Foxbase traffic, performance across the Persoft bridge slows by 23 to 25 percent (i.e., compared with the No Bridge times).

IPX traffic slows the Hughes significantly (possibly because of high packet drop rates; see the Dropped Packets section). With no traffic and with Lotus and Foxbase traffic, the Hughes outperforms the CrossComm and competes with the Netronix.

Xcopy

A 5M-byte directory tree (130 files in 13 directories) is copied from the primary workstation to the server (and vice versa) using the DOS Xcopy command with the /s parameter. Both tests run without traffic, with Lotus 1-2-3 and Foxbase traffic, and with intense IPX traffic on both segments.

Figure 11a.
Xcopy to Server

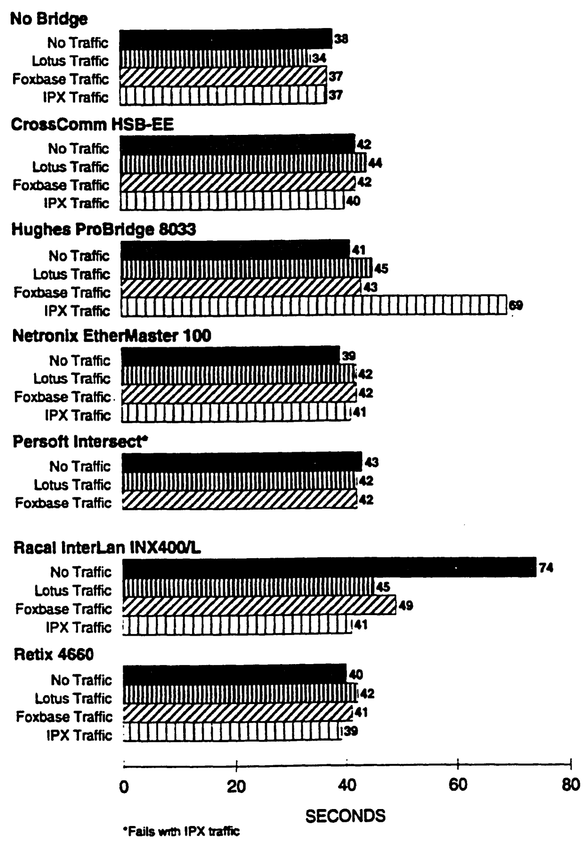

Figure 11b.
Xcopy from Server

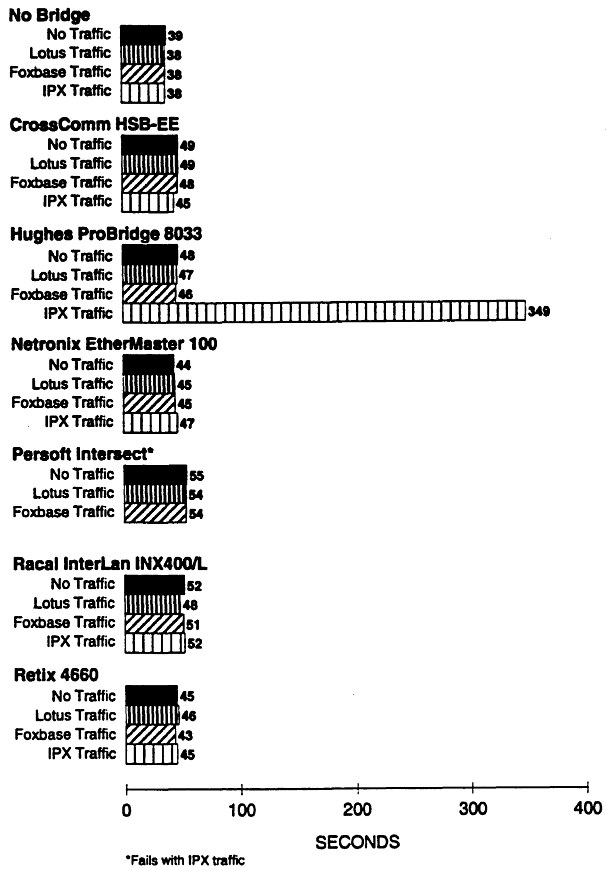

Analysis

The large volume of data moving between the workstation and server produces more contrast between the bridge times and the baseline No Bridge times than in some of the other tests.

As in the other tests, the Retix returns superior performance with all types of traffic, and the CrossComm and Netronix perform consistently well overall. The Racal is at its slowest with no traffic and at its fastest with intense IPX traffic. The Persoft performs well except with IPX traffic, which causes it to fail. IPX traffic slows the Hughes somewhat when data moves to the server (data and traffic are moving in the same direction across the bridge) and slows precipitously when data moves from the server to the workstation (data and traffic are moving in opposing directions across the bridge).

Foxbase+/LAN

Foxbase+/LAN executes a series of database operations. The tests run without traffic, with Lotus 1-2-3 and Foxbase traffic, and with intense IPX traffic on both segments.

Foxbase+/LAN database operations include:

1. Writes 51,100 records to three different files, testing the speed of writing records.
2. Indexes the Account (50,000 records), Teller (1,000 records), Transact (1,000 records), and Branch (100 records) files.
3. Copies the first 3 of every 10 Account records to three temporary files.
4. Deletes the copied records (Test 3) from the Account file.
5. Appends the temporary files (Test 3) to the Account file.
6. Packs and reindexes the Account file.
7. Processes 1,000 transactions, each updating a record in the Account, Teller, and Branch tables and writing a History record.
8. Creates four indexes on the 1,000-record History file.
9. Repeats Test 7 and updates the History file indexes.
10. Appends blocks of 1,000-records to the History file.
11. Repeats Test 9 with the 10,000-record indexed History file.
12. Selects 1,000 Account records in indexed order and prints the file to disk.
13. Selects 1,000 Account records on an unindexed field and prints a report to disk.
14. Groups and subtotals the 1,000-record Teller file and prints the report to disk.
15. Joins the Teller and Branch files and prints the report to disk.
16. Joins the History, Branch, Teller, and Account files (1,000, 100, 1,000, and 50,000 records), sorts the History file, and prints the report to disk.

Analysis

The bridges exhibit rather homogeneous performance. Times are especially close for operations of relatively short duration; differences are more perceptible for operations of longer duration (notably Tests 4, 6, 10, and 11), possibly because of the cumulative effect of the bridges' different internal

Table 1. Foxbase+/LAN Tests

No Traffic (times in seconds)

	No Bridge	CrossComm HSB-EE	Hughes ProBridge 8033	Netronix EtherMaster 100	Persoft Intersect	Racal InterLan INX400/L	Retix 4660
Test 1	47	49	49	48	51	49	48
Test 2	15	18	16	17	19	17	16
Test 3	18	25	24	23	28	24	22
Test 4	25	36	35	32	42	34	31
Test 5	8	12	12	10	13	11	10
Test 6	22	29	28	27	33	28	26
Test 7	19	21	21	20	21	21	20
Test 8	4	6	5	5	6	6	5
Test 9	30	36	36	35	39	36	33
Test 10	55	68	66	63	73	65	61
Test 11	41	55	52	49	57	52	46
Test 12	7	7	7	7	7	7	7
Test 13	13	15	16	14	17	15	14
Test 14	2	2	2	2	1	1	2
Test 15	2	2	2	2	1	2	2
Test 16	20	22	21	22	24	22	21

Lotus Traffic (times in seconds)

	No Bridge	CrossComm HSB-EE	Hughes ProBridge 8033	Netronix EtherMaster 100	Persoft Intersect	Racal InterLan INX400/L	Retix 4660
Test 1	47	50	50	49	51	49	48
Test 2	15	18	17	17	20	17	16
Test 3	18	25	23	23	28	24	23
Test 4	25	36	36	32	43	34	31
Test 5	9	13	11	11	14	12	10
Test 6	22	29	28	27	33	29	26
Test 7	19	21	20	20	21	20	20
Test 8	4	6	6	6	7	6	6
Test 9	30	36	36	34	39	36	32
Test 10	56	67	65	64	75	65	63
Test 11	41	55	51	49	58	53	47
Test 12	7	7	7	7	8	7	7
Test 13	13	15	16	15	17	15	14
Test 14	1	2	1	1	2	2	2
Test 15	2	2	2	2	2	2	2
Test 16	20	23	22	22	23	22	21

Foxbase Traffic (times in seconds)

	No Bridge	CrossComm HSB-EE	Hughes ProBridge 8033	Netronix EtherMaster 100	Persoft Intersect	Racal InterLan INX400/L	Retix 4660
Test 1	47	50	49	49	51	49	48
Test 2	14	17	17	16	19	17	17
Test 3	19	25	24	23	29	24	32
Test 4	25	36	34	33	42	34	31
Test 5	8	11	11	11	14	11	11
Test 6	22	28	27	26	33	28	26
Test 7	19	21	21	20	21	21	20
Test 8	5	6	5	5	7	5	5
Test 9	30	35	36	35	39	36	33
Test 10	55	67	65	63	74	65	62
Test 11	42	54	53	49	58	52	48
Test 12	7	7	7	7	8	8	6
Test 13	13	16	15	15	17	15	14
Test 14	2	2	1	2	2	2	2
Test 15	1	2	2	2	2	2	2
Test 16	21	22	22	22	23	23	21

IPX Traffic (times in seconds)

	No Bridge	CrossComm HSB-EE	Hughes ProBridge 8033	Netronix EtherMaster 100	Persoft Intersect	Racal InterLan INX400/L	Retix 4660
Test 1	47	48	63	50	NA	49	48
Test 2	15	17	73	16	NA	17	17
Test 3	18	23	176	23	NA	24	22
Test 4	25	32	197	34	NA	35	32
Test 5	8	10	49	10	NA	11	10
Test 6	22	27	163	28	NA	29	27
Test 7	19	20	26	20	NA	20	20
Test 8	5	5	21	5	NA	6	5
Test 9	30	33	59	35	NA	36	33
Test 10	55	63	278	64	NA	65	63
Test 11	42	49	94	49	NA	53	48
Test 12	7	7	13	7	NA	7	6
Test 13	13	14	78	15	NA	15	15
Test 14	1	2	5	2	NA	2	2
Test 15	2	2	6	2	NA	2	2
Test 16	20	21	40	21	NA	22	21

NA—Not applicable.

Figure 12a.
Internal Latency, No Traffic

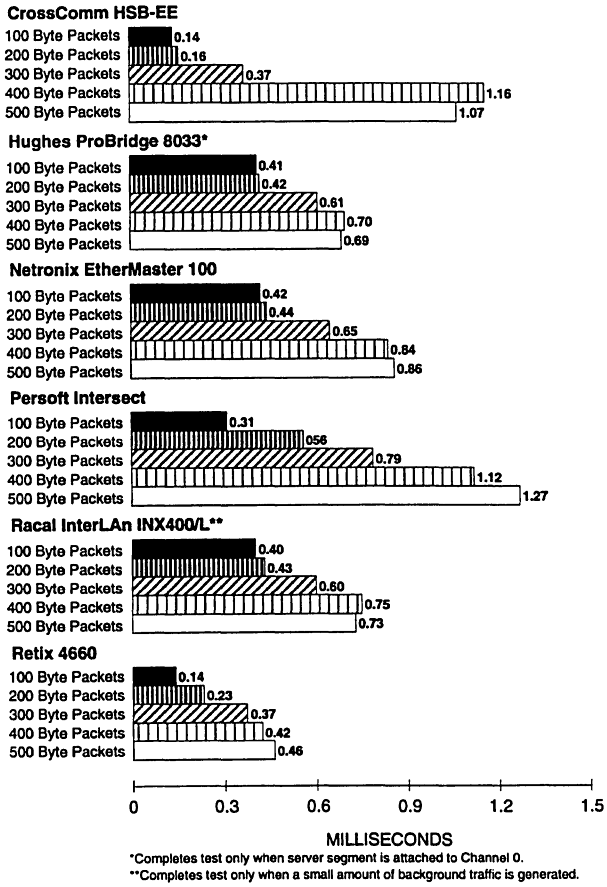

latencies. The CrossComm is somewhat slower than the Netronix, Racal, and Hughes; CrossComm's internal latency without traffic is distinctly nonlinear with respect to variations in packet size.

Internal Latency

Ethernet frames in five sizes are generated and transmitted back and forth across the bridge to measure bridge processing overhead associated with IPX packets of various sizes. NSTL's custom IPX send and receive programs use low-level IPX system calls and specific IPX sockets to generate and send each frame. After sending each frame, the IPX send program waits for an acknowledgment frame from the receive program before sending another frame. If an acknowledgment packet is not received within three seconds, a dropped packet is logged and the send program transmits another frame. Frames transmitted by the send program

Figure 12b.
Internal Latency, IPX Traffic

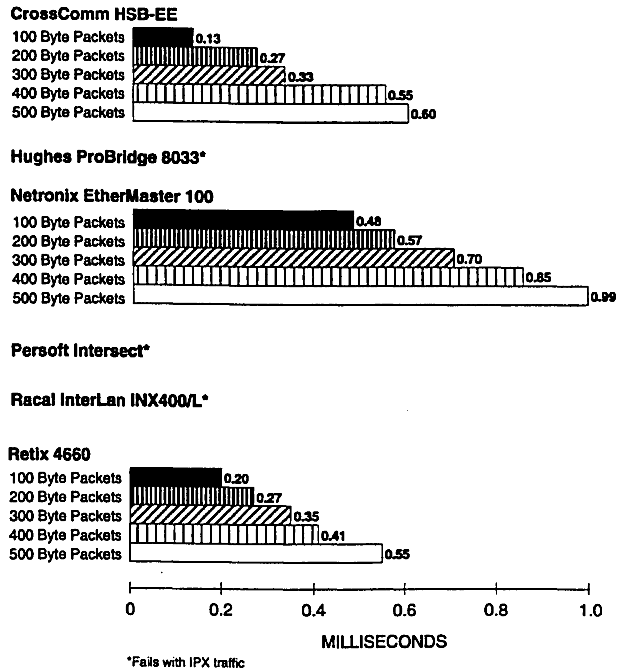

are destination specific, and those transmitted by the receive program are broadcast. Latency refers to the delay caused by the bridge forwarding packets to another segment rather than filtering the packets within a single network segment. Gross latency calculations are based only on runs completed without dropped packets. Net latencies for each bridge subtract the latency values for transmissions without a bridge, effectively normalizing for frame processing overhead not related to the bridge. Transmissions at all packet sizes are repeated with intense IPX traffic on both network segments.

Analysis

Latency measurements with no background traffic reveal a somewhat linear relationship between latency and packet size. The CrossComm exhibits a large latency increase at the 400-byte packet interval, indicating that the linear relationship between latency and packet size may depend to some extent on traffic intensity. Latency-packet size relationships become noticeably more linear with intense IPX background traffic. In fact, the CrossComm produces smaller latencies overall with IPX traffic than without.

The Racal performs at its best with traffic and fails completely when no traffic is present on the network. Small amounts of broadcast background traffic enable the Racal to complete the test. The Sniffer analysis reveals that the Racal requires small groups of packets forwarded at regular intervals which coincide with the regular transmission of "hello" messages transmitted by the bridge (a characteristic of spanning-tree bridges). The Racal bridge does not successfully transmit the number of frames required to complete the latency test, but constant low-intensity broadcast traffic generated by the Sniffer causes the Racal to correctly forward all the frames. Although contrary to the behavior of most bridges (which forward packets expeditiously when no traffic is present and slowly with heavy traffic), networks requiring a bridge only rarely experience periods with absolutely no background traffic. Small amounts of network traffic during normal business hours will mask the problem, making it difficult to diagnose.

The Persoft takes the longest to process and forward packets. According to the manufacturer, the Persoft should show performance gains in a more powerful system.

Dropped Packets

Ethernet packets are forwarded from one network segment to the other. NSTL's custom IPX Blaster program transmits 65,535 small packets across the bridge at 7,281 frames per second. A Network General Sniffer running Ethernet Monitor Version 3.0 counts the frames received on the destination segment. This test runs without background traffic.

Some packets are dropped or lost during normal Ethernet network operation. High-level protocols call for retransmission of dropped packets, and minimal drop rates produce negligible retransmission delays. With higher drop rates, the cumulative delay from retransmissions can noticeably degrade network performance.

Analysis

The dropped packet test provides a gross measurement of a bridge's capability to handle high traffic volumes. Because the uniformly small frames require minimal processing overhead, the test is biased in favor of bridges with short latencies for small packets and in favor of bridges with efficient frame buffering schemes.

The Hughes drops over 8 percent of its packets; this high drop rate may explain the Hughes bridge's poor benchmark results with intense IPX traffic and its inability to complete the latency test with IPX traffic. According to the manufacturer's specifications, the Hughes should sustain a forwarding rate of 13,000 packets per second with learning enabled, but measurable performance degradations are observed at lower forwarding rates.

Persoft does not publish specific filtering or forwarding rates for its Intersect bridge, but it does claim more powerful performance in a more powerful system. As configured for testing, the Persoft bridge shows a noticeable performance loss when confronted with forwarding 7,281 frames per second, fails completely with IPX traffic, exhibits the longest latencies for all packet sizes, and performs the slowest overall. A more powerful host may improve the Persoft's performance.

Table 2. Dropped Packets

Vendor/Product	Averaged Dropped Packet Rate (%)
CrossComm HSB-EE	0.0
Hughes ProBridge 8033	8.1
Netronix EtherMaster 100	0.0
Persoft Intersect	26.9
Racal InterLan INX400/L	0.0
Retix 4660	0.0

Vendors

CrossComm Corp.
140 Locke Drive
Marlborough, MA 01752 (508) 481-4060

Hughes LAN Systems
1225 Charleston Road
Mountain View, CA 94043 (800) 395-5267

Netronix, Inc.
1372 N. McDowell Boulevard
Petaluma, CA 94954 (707) 762-2703, (800) 762-2703

Persoft, Inc.
465 Science Drive
Madison, WI 53711 (608) 273-6000

Racal InterLan
155 Swanson Road
Boxborough, MA 01719 (508) 263-9929, (800) LAN-TALK

Retix
2644 30th Street
Santa Monica, CA 90405 (213) 399-2200

Characteristics

Table 3. Local Ethernet Bridge Characteristics

	Filtering/Forwarding Rates	Address Table Size	Management Protocols	Memory (bytes)	Processor	Warranty and Support
CrossComm HSB-EE	Filtering: 15,000 pps; Forwarding: 14,800 pps	Transient: 40,000; Permanent: 128	SNMP	Address Table RAM: 128K; Frame Buffer RAM: 128K ultra-high-speed RAM	80376	3 months, parts, labor, and return shipment; extended warranty available ($500); technical support ($125 per hour; free with support contract); telephone support
Hughes Pro-Bridge 8033	Filtering: 25,000 pps; forwarding: 13,500 pps	Transient: 3,000; Permanent: 800	SNMP	Address Table RAM: 96K; Frame Buffer RAM: 64K	80386SX	1 year, parts, labor, and return shipment; extended warranty available ($289); telephone support and updates free for 3 months; $700 per year thereafter
Netronix Ether-Master 100	Filtering: 25,000 pps; Forwarding: 14,000 pps	Transient: 8,000; Permanent: 1,000	SNMP	Address Table RAM: 114K; Frame Buffer RAM: 256K	80376	1 year, parts, labor, and return shipment; telephone support
Persoft Intersect	Filtering: 22,000 pps; Forwarding: 11,000 pps	8,132	SNMP support planned for next release	Host dependent	Host dependent	5 years, parts, labor, and return shipment; 30-day money-back guarantee; telephone support
Racal InterLan INX400/L	Filtering: 23,600 pps; Forwarding: 11,000 pps	Transient: 8,000; Permanent: 50	SNMP	Address Table RAM: 700K; Frame Buffer RAM: 91K	80286	1 year, parts, labor, and return shipment; telephone support
Retix 4660	Filtering: 29,000 pps; Forwarding: 13,650 pps	Transient: 8,192; Permanent: 1,000	SNMP, CMIP, MAP, TOP	Address Table RAM: 96K; Frame Buffer RAM: 200K	68020	1 year, parts, labor, and return shipment; warranty repairs through NCR; telephone support

Table 4. Local Ethernet Bridge Features

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix EtherMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
Physical Characteristics	0						
Rack-mount option	0	▲	▲	▲	NA	▲	▲
Weight (lb.)	0	10	6.5	18	NA	NA	13
Height (in.)	0	2.6	7	4.5	NA	2.7	3.3
Depth (in.)	0	12.5	3.4	15.5	NA	17	12.5
Length (in.)	0	15.1	14.5	16.5	NA	13.8	17
Electrical/Environmental	0						
Power requirements (Watts)	0	60	65	150	NA	232	50
Operating humidity (%), noncondensing	0	5-90	10-90	5-95	NA	10-90	5-95
Operating temperature (degrees C)	0	0-40	10-40	10-55	NA	0-50	0-50
FCC class (A or B)	0	A	A	B	NA	A	A
Switchable 110/220 (115/230) V AC	0	▲	—	▲	NA	▲	▲
Cooling fan	0	▲	—	▲	NA	▲	▲
Power switch	0	▲	—	—	NA	▲	▲
Reset button	0	▲	▲	▲	NA	—	—
Processor	0						
Processor type	0	80376	386SX	80376	Host	80286	68020
Processor speed (MHz)	0	16	16	16	Host	10	20
Embedded controller	0	▲	▲	▲	—	—	—
Memory	0						
Address table size (KB)	0	128	96	114	NA	700	96
Frame buffer size (KB)	0	128**	64	256	NA	91	200
Physical Connections (Quantity)	1						
DB15	1	2	2	2	2	2	2
BNC	1	2	2	2	2	2	2
Fiber	1	0	0	0	0	0	2
RS-232-C	1	1	1	0	0	1	1
Media	1						
Thick coaxial	1	▲	▲	▲	▲	▲	▲
Thin coaxial	1	▲	▲	▲	▲	▲	▲
Twisted pair	1	—	—	—	▲	▲	▲
Fiber	1	—	—	—	—	—	▲
Diagnostics	1						
Power-up diagnostics	5	▲	▲	▲	▲	▲	▲
LED/LCD/Tone indicators	3	▲	▲	▲	▲	▲	▲
Console command line diagnostic	3	▲	▲	—	—	▲	▲
Downloadable diagnostics	1	▲	—	—	—	—	▲
Front-Panel LED/LCD Indicators	1						
Power (AC)	1	▲	▲	▲	—	▲	▲

Table 4. Local Ethernet Bridge Features (Continued)

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix Eth-erMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
Alarm (12-Volt power failure)	1	—	—	▲	—	—	▲
Self-test/diagnostic	1	▲	▲	▲	—	▲	▲
Console	1	▲	▲	—	—	—	▲
Network A collision	1	▲	—	▲	—	▲	▲
Network A transmit	1	▲	▲	▲	—	▲	▲
Network A receive	1	▲	—	▲	—	▲	▲
Network B collision	1	▲	—	▲	—	▲	▲
Network B transmit	1	▲	▲	▲	—	▲	▲
Network B receive	1	▲	—	▲	—	▲	▲
Administrative Software Connection	1						
RS-232 (TTY) attachment	1	▲	▲	—	▲	▲	▲
Network attached	1	▲	▲	▲	—	▲	▲
Modem support	1	▲	▲	—	—	▲	▲
Network Management Support	2						
Bridge management software	5	▲	—	▲	—	▲	▲
Bridge console commands	5	▲	▲	—	—	▲	▲
SNMP (Simple Network Management Protocol)	5	▲	▲	▲	—	▲	▲
CMIP (Common Management Information Protocol)	1	—	—	—	—	—	▲
CMOT (Common Management Protocol over TCP/IP)	1	—	—	—	—	—	—
MAP (Manufacturing Automation Protocol)	1	—	—	—	—	—	▲
TOP (Technical Office Protocol)	1	—	—	—	—	—	▲
IEEE 802.3 Standards	1						
10BASE5 (standard/thick Ethernet)	2	▲	▲	▲	▲	▲	▲
10BASE2 (thin Ethernet)	2	▲	▲	▲	▲	▲	▲
10BASE-T (unshielded twisted-pair Ethernet)	2	—	—	—	▲	▲	▲
FOIRL (Fiber Optic Internetwork Repeater Link)	2	—	—	—	—	—	▲
High-Level Protocol Transparency	1						
TCP/IP	2	▲	▲	▲	▲	▲	▲
DECnet	2	▲	▲	▲	▲	▲	▲
LAT	2	▲	▲	▲	▲	▲	▲
XNS	2	▲	▲	▲	▲	▲	▲
ISO	2	▲	▲	▲	▲	▲	▲
Operating System Transparency	1						
Novell NetWare	2	▲	▲	▲	▲	▲	▲

Table 4. Local Ethernet Bridge Features (Continued)

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix Eth-erMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
3Com	2	▲	▲	▲	▲	▲	▲
Banyan	2	▲	▲	▲	▲	▲	▲
IEEE 802.3 compatibles	2	▲	▲	▲	▲	▲	▲
Performance (Packets per Second)*	0						
Filtering	1	15K	25K	25K	22K	23.6K	29K
Forwarding, spanning-tree/learning enabled	1	14,880	13K	12.5K	11K	11K	13,650
Forwarding, spanning-tree/learning disabled	1	14,880	13,500	14K	NA	NA	13,650
Throughput (M bps)	1	9.8	10	9.8	10	NA	10
Source Address Database	3						
Database Functions							
Automatic address learning	5	▲	▲	▲	▲	▲	▲
Maximum address table entries	1	40K	3K	8K	8,132	8K	8,192
Maximum addresses in nonvolatile RAM	2	128	800	1K	NA	50	4K
Filter/forward based on address range	1	—	—	▲	—	—	▲
Filter/forward based on database entries	5	▲	▲	▲	▲	▲	▲
Automatic aging	2	—	▲	▲	—	▲	▲
Accelerated aging upon spanning-tree change	1	—	▲	▲	—	▲	—
Manually add records	2	▲	▲	▲	—	▲	▲
Manually delete records	2	▲	▲	▲	—	▲	▲
Exempt record from aging	1	—	▲	—	—	▲	—
Lock record characteristics	1	▲	▲	—	—	—	▲
Auto correct subnet assignments	1	—	▲	—	—	▲	▲
Save records in nonvolatile RAM	2	▲	▲	▲	—	▲	▲
Assign symbolic names to records	1	▲	▲	—	—	—	—
Add address to range table	1	—	—	▲	—	—	▲
Delete address from range table	1	—	—	▲	—	—	▲
Save/retrieve range table configurations	1	—	—	▲	—	—	▲
Display Options							
All records	1	▲	▲	▲	—	▲	▲
Restricted records	1	▲	▲	—	—	▲	▲
Static records	1	▲	▲	▲	—	▲	▲
Records in nonvolatile RAM	1	▲	▲	▲	—	▲	▲

Table 4. Local Ethernet Bridge Features (Continued)

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix Eth-erMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
Settable Options							
Aging time interval	1	—	▲	▲	—	▲	—
Symbolic names on/off	1	▲	▲	—	—	—	—
Record disposition	2	—	▲	▲	—	▲	▲
Range table entry limits (upper/lower)	1	—	—	▲	—	—	▲
Range table entry disposition	1	—	—	▲	—	—	▲
Protocol/Operating System Filtering	2						
AppleTalk	1	—	▲	▲	—	▲	▲
ARP	1	▲	▲	▲	—	▲	▲
Bridge ID	1	—	▲	▲	—	NA	▲
IP	1	▲	▲	▲	—	▲	▲
DECnet	1	▲	▲	▲	—	▲	▲
LAT	1	▲	▲	▲	—	▲	▲
NetWare	4	▲	▲	▲	—	▲	▲
REVARP	1	▲	▲	▲	—	▲	▲
TCP/IP	1	▲	▲	▲	—	—	▲
TOP	1	—	▲	—	—	▲	▲
V2	1	—	▲	▲	—	NA	▲
XNS	1	▲	▲	▲	—	▲	▲
User-defined protocol	1	—	▲	—	—	▲	—
Forwarding	1						
Unknown destination address	5	▲	▲	▲	▲	▲	▲
Broadcast frames	5	▲	▲	▲	▲	▲	▲
Group address	1	▲	▲	▲	—	—	▲
Individual address	5	▲	▲	▲	▲	▲	▲
Routing Methods	1						
Source address	1	▲	▲	▲	▲	▲	▲
802.1 spanning-tree protocol	1	▲	▲	▲	▲	▲	▲
Source routing transparency	1	—	—	—	—	—	—
Automatic cable loop avoidance	1	▲	▲	▲	▲	▲	▲
Administration/Configuration Functions	3						
Management software included	1	▲	▲	▲	—	▲	▲
Help function	1	▲	▲	▲	—	▲	▲
Monitor multiple bridges	1	▲	—	▲	—	▲	▲
Menu-driven management	1	▲	—	▲	—	▲	▲
Save/retrieve profile configuration	1	▲	▲	▲	—	—	▲
Password access to Management	1	▲	▲	—	—	▲	▲

Table 4. Local Ethernet Bridge Features (Continued)

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix Eth-erMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
Multilevel password access to Management	1	—	—	—	—	▲	▲
Modify console serial connection settings	1	▲	▲	—	—	▲	▲
Console command line Help	1	—	▲	▲	—	—	▲
Bridge off-/on-line	1	▲	▲	▲	—	▲	▲
Network device listing by manufacturer	1	—	▲	—	—	—	▲
Automatic update of reset log	1	—	▲	—	—	—	▲
Always forward specific types (IP, ARP, etc.)	1	▲	▲	▲	—	▲	▲
Software reset with updated configuration	1	▲	▲	▲	—	▲	▲
Hardware reset with updated configuration	1	—	▲	▲	▲	▲	▲
Software reset with default configuration	1	▲	▲	▲	—	—	▲
Hardware reset with default configuration	1	—	▲	▲	—	▲	▲
Report Network Management Center IP address	1	—	▲	—	—	—	▲
Change SNMP password	1	—	▲	▲	—	—	▲
Console control over multiple bridges	1	▲	—	—	—	—	▲
Save settings to nonvolatile RAM	1	▲	▲	▲	—	▲	▲
Modular software cartridge	1	—	▲	—	—	▲	▲
Reset all statistics	1	▲	▲	▲	—	▲	▲
Initialize start-up/diagnostic cycle	1	▲	▲	—	—	▲	▲
Logout of Management session	1	—	▲	▲	—	▲	▲
Reject specific frame types	1	▲	▲	▲	—	▲	▲
Accept specific frame types	1	▲	▲	▲	—	▲	▲
Forwarding of multicast frames on/off	1	▲	▲	▲	—	▲	▲
Integrated text editor	1	—	—	▲	—	—	▲
Administrative Display Functions							
Bridge configuration/status	1	▲	▲	▲	▲	▲	▲
Parameter-specific status	1	▲	▲	▲	—	▲	▲
Bridge name	1	▲	▲	▲	▲	▲	▲
Reset log	1	▲	▲	—	—	—	▲
Name of downloadable software file	1	▲	—	▲	—	▲	▲
Console event reporting level	1	—	▲	—	—	▲	▲

Table 4. Local Ethernet Bridge Features (Continued)

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix Eth-erMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
Bridge operation mode	1	▲	▲	▲	▲	▲	▲
Names of attached segments	1	—	▲	—	▲	▲	▲
Login message	1	—	—	—	—	▲	▲
Administrative Display Functions, Channel 0							
Bridge Ethernet address	1	▲	▲	▲	▲	▲	▲
All bridge connections	1	▲	▲	—	—	—	▲
Packets transmitted	1	▲	▲	▲	—	▲	▲
Packets received	1	▲	▲	▲	—	▲	▲
CRC and alignment errors	1	▲	▲	▲	—	▲	▲
No-resource errors	1	▲	▲	▲	—	▲	▲
Oversize packets	1	▲	▲	—	—	▲	▲
Undersize packets	1	▲	▲	▲	—	▲	▲
Packet transmission failures	1	▲	▲	▲	—	▲	▲
Packet transmission collisions	1	▲	▲	▲	—	▲	▲
Packets queued for transmission	1	—	▲	—	—	▲	▲
SNMP packets broadcast	1	▲	▲	—	—	—	▲
Valid SNMP packets transmitted	1	▲	▲	—	—	—	▲
Administrative Display Functions, Channel 1							
Packets transmitted	1	▲	▲	▲	—	▲	▲
Packets received	1	▲	▲	▲	—	▲	▲
CRC and alignment errors	1	▲	▲	▲	—	▲	▲
No-resource errors	1	—	▲	▲	—	▲	▲
Oversize packets	1	▲	▲	▲	—	▲	▲
Undersize packets	1	▲	▲	▲	—	▲	▲
Packet transmission failures	1	▲	▲	▲	—	▲	▲
Packet transmission collisions	1	▲	▲	▲	—	▲	▲
Packets queued for transmission	1	▲	▲	—	—	▲	▲
SNMP packets broadcast	1	▲	▲	—	—	—	▲
Valid SNMP packets transmitted	1	▲	▲	—	—	—	▲
Settable Administrative Functions							
Packet protocols to be blocked	1	▲	▲	▲	—	▲	▲
Console command line prompt	1	—	▲	—	—	▲	▲
Bridge name	1	▲	▲	▲	—	▲	▲
Attached segment names	1	—	▲	—	—	▲	▲
Bridge IP address	1	▲	▲	—	—	▲	▲

Table 4. Local Ethernet Bridge Features (Continued)

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix EtherMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
IP net mask for IP routing	1	—	▲	—	—	▲	▲
IP router address	1	—	▲	—	—	▲	▲
Origin filtering on/off	1	—	▲	▲	—	—	▲
Standard/extended packet sizes	1	—	▲	—	—	▲	▲
Restrict mode on/off	1	—	▲	—	—	—	▲
Name downloadable bridge software file	1	▲	▲	▲	—	▲	▲
Console event reporting level	1	—	▲	—	—	▲	▲
Forward operating mode	1	▲	▲	▲	—	▲	▲
Filter operating mode	1	▲	▲	▲	—	▲	▲
Pass all mode	1	▲	▲	▲	—	▲	▲
Filter/forward greater than "size"	1	—	▲	—	—	▲	▲
Filter/forward smaller than "size"	1	—	▲	—	—	▲	▲
Login message	1	—	—	—	—	▲	▲
Statistics event threshold options	1	—	—	—	—	—	▲
Statistics gathering options	1	—	▲	▲	—	▲	▲
Bridge system password	1	▲	▲	▲	—	▲	▲
Bridge user (view only) password	1	▲	—	—	—	▲	▲
Source address to filter only	1	▲	▲	▲	—	▲	▲
Source address to forward only	1	▲	—	▲	—	▲	▲
Time-out limit for transmission reset	1	—	▲	—	—	—	▲
Time-out limit for receive reset	1	—	▲	—	—	—	▲
Forwarding delay	1	—	▲	▲	—	▲	▲
Administrative Set Functions, Channel 0							
Source address learning mode on/off	1	▲	▲	—	—	—	▲
Source address record aging on/off	1	▲	▲	—	—	▲	▲
Port 1 state	1	▲	▲	▲	—	—	▲
Port 1 priority level	1	—	▲	▲	—	—	▲
Administrative Set Functions, Channel 1							
Source address learning mode on/off	1	▲	▲	—	—	—	▲
Source address record aging on/off	1	▲	▲	—	—	▲	▲
Port 2 state	1	▲	▲	▲	—	—	▲
Port 2 priority level	1	—	▲	▲	—	—	▲

Table 4. Local Ethernet Bridge Features (Continued)

	Weight	CrossComm HSB-EE	Hughes Pro-Bridge 8033	Netronix Eth-erMaster 100	Persoft Intersect	Racal Inter-Lan INX400/L	Retix 4660
Administrative Flag Settings							
Error severity flags	1	—	—	—	—	—	—
Forwarding flag, unknown group addresses	1	—	▲	—	—	—	▲
Forwarding flag, unknown individual addresses	1	—	▲	—	—	—	▲
Origin filter learning flag, source addresses	1	—	▲	—	—	▲	▲
Restricted learning flag, destination addresses	1	—	▲	—	—	—	▲
Static learning flag	1	—	▲	▲	—	—	▲
Administrative Spanning-Tree Settings							
Bridge standby mode (backup)	1	—	▲	▲	—	▲	▲
Root priority level	1	▲	▲	▲	—	▲	▲
Hello time	1	▲	▲	▲	—	▲	▲
Bad hello limit	1	▲	—	—	—	▲	▲
New tree time	1	▲	—	—	—	▲	▲
Hello maximum age	1	▲	▲	▲	—	▲	▲
Secure mode	1	—	—	▲	—	—	▲
Enable/disable spanning-tree protocol	1	—	▲	▲	—	▲	▲
Path cost per segment	1	▲	▲	▲	—	▲	▲

NA—Not applicable.
**64-byte packets per manufacturer's specification.*
***Ultrahigh-speed RAM.*

Glossary of Terms

Selected Definitions

Source Address Database Functions

Auto-Correct Subnet Assignments: A bridge tracks the address and Ethernet segment location (subnet assignment) of each node (workstation with an Ethernet card having a unique Ethernet address) on each segment connected to the bridge. When a node is disconnected from one segment and connected to the other, a bridge with automatic correction updates its internal address record to reflect the node's new location. The correction goes into effect as soon as the workstation transmits a frame from its new location.

Assign Symbolic Names to Records: A symbolic name can be assigned to an Ethernet address to simplify references; for example, "Super 486 Workstation" in place of a 12-character hexadecimal string such as 0000C0EA2B1C.

Range Table Entries: This term refers to a database of Ethernet address ranges referenced by the bridge. The bridge either filters or forwards a frame depending on whether it is within a specified range.

Display Options

Restricted Records: The bridge can be set up to selectively restrict forwarding of packets destined for specific addresses on a segment.

Static Records: Address records in the source address database can be designated as "static," and these records will not be changed by aging and learning of addresses.

Settable Options

Aging Time Interval: The administrator can specify a time interval after which a nonstatic address record will be purged from the address database unless the bridge has received a packet from it.

Record Disposition: The administrator can specify actions to be taken by the bridge upon receiving a packet from a particular address—always forward the packet to the other segment, forward the packet to port 1 if received on port 2, forward to port 2 if received on port 1, and never forward the packet.

Forwarding

Group Address: Frames with group or multicast addresses designate more than one station, but, unlike broadcast frames, the frames are not received by all stations.

Administration/Configuration Functions

Report Network Management Center IP Address: Simple Network Management Protocol (SNMP) uses IP addresses to establish communication between centralized network management stations and devices such as bridges on the network.

Forwarding of Multicast Frames on/off: The administrator can optionally designate that all packets destined for group or multicast addresses be automatically forwarded by the bridge.

Administrative Display Functions

Console Event Reporting Level: The administrator can specify that a console attached to the bridge display local event messages only or display all event messages on the network. Local event messages are those generated by the bridge.

Bridge Operation Modes: The administrator can set and display the current bridge operation mode; modes typically include Backup, Learn, Forward, Pass All, and some kind of Restriction-by-Packet-Size.

Administrative Display Functions, Channel 0

CRC and Alignment Errors: The bridge can report errors in receiver synchronization, invalid source or destination address, and invalid data fields according to a Frame Check Sum (FCS).

No-Resource Errors: The bridge can report the number of packets not received owing to insufficient frame buffer space.

SNMP Packets Broadcast: The bridge can report the number of valid SNMP packets transmitted on either channel.

Settable Administrative Functions

IP Net Mask for IP Routing: The bridge can determine whether a destination address is part of the local segment attached to the bridge.

IP Router Address: The administrator can designate an IP router address. When the bridge determines that a destination address is located outside the local subnet, packets can be passed to the IP router.

Standard/Extended Packet Sizes: The administrator can designate that the bridge accept the standard IEEE 802.3 Ethernet packet size (60 to 1,536 bytes) or an extended range of packet sizes.

Statistics Event Threshold Options: The bridge can report events based on threshold levels of activity (e.g., maximum value of errors) or rates (e.g., maximum frequency of errors).

Time-Out Limit for Reset of the Transmitter or Receiver Chip: The administrator can specify a time interval after which the bridge's receiver and transmitter chips reset when no transmission occurs due to cable failure or disconnection.

Administrative Spanning-Tree Settings

Root Priority Level: The administrator can use a spanning tree to determine which bridge is the root bridge.

Hello Time: A spanning-tree setting specifies the interval at which a bridge must send a message identifying itself as the root bridge of the spanning tree.

Bad Hello Limit: A bad hello message is one transmitted by a bridge other than the root bridge; when the limit for bad hellos is exceeded, a warning event message may be issued.

New Tree Time: A spanning-tree setting specifies an interval after which a spanning-tree bridge discards its current configuration information and reconfigures the spanning tree.

Hello Maximum Age: A spanning-tree setting specifies the maximum time period that a hello message generated by a bridge can remain on the network before it "times out" and is discarded.

Secure Mode: The source address database can be secured such that the bridge will forward only packets from addresses manually entered (versus automatically learned) into the database.

Path Cost per Segment: Cost values can be assigned to the transmission of packets across a bridge. Costs are additive as more bridges are traversed. The spanning-tree algorithm uses path costs to determine the optimum path.

Product Prices

		Purchase Price ($)
Local Ethernet Bridges		
	CrossComm HSB-EE	4,350
	Hughes ProBridge 8033	2,995
	Netronix EtherMaster 100	3,400
	Persoft Intersect	1,495
	Racal InterLan INX400/L	3,195
	Retix 4660	3,750

■

Remote Ethernet Bridges

In this report:

A Report from NSTL

Synopsis

Focus

This report discusses applciations for remote Ethernet bridges and evaluates six products' performance, features, and usability.

Products Tested

HP 28674A
Hewlett-Packard Co.

3000 Series HCB
Magnalink Communications Corp.

REB-20
RAD Network Devices

High Performance 4880
Retix

NETBuilder IB/3000
3Com Corp.

TransLAN 335
Vitalink Communications Corp.

Product Recommendations

- Magnalink 3000 Series HCB
- Hewlett-Packard HP 28674A
- Retix High Performance 4880

Source

Based on data generated by tests designed and conducted by National Software Testing Laboratories, Inc. (NSTL), a division of Datapro Information Services, Plymouth Meeting, PA 19462. Telephone (800) 223-7093.

Ratings Key
(On a scale of 0 to 10)
Ratings
● 7.0 - 10.0
○ 5.0 - 6.9
◍ under 5.0
★ Recommended

Overall Evaluation	Product Name	Performance	Features	Usability	Price
8.0	Magnalink 3000 Series HCB	●	●	●	$9,400 ★
7.7	Hewlett-Packard HP 28674A	●	○	●	$3,999 ★
7.7	Retix High Performance 4880	●	○	●	$8,400 ★
7.4	3Com NETBuilder IB/3000	●	●	●	$7,495
6.7	Vitalink TransLAN 335	○	●	●	$11,500
5.9	RAD REB-20	◍	○	●	$6,950

Overview

Remote Ethernet bridges connect multiple LANs across large geographic distances. Connectivity of this nature enables sharing of important resources among workgroups at physically separate locations. Such "enterprise-wide" networks become more and more feasible as wide area communications equipment and data links become more powerful, less expensive, and more available.

With the inexorable transition of North America's analog Public Switched Telephone Network to a completely digital Integrated Services Digital Network (ISDN), connectivity over large distances using digital communications will become more commonplace. Many internetworking products already address some of these needs, and as their popularity increases, their capabilities will likely increase as well.

Evaluation Criteria

NSTL evaluated IEEE 802.3-compatible bridges capable of bridging two remote Ethernet network segments at the Media Access Control (MAC) layer of the ISO model. MAC-layer bridges filter or forward packets after examining MAC-layer address information contained within the first 12 bytes of the packet. All the bridges support at least one wide area connection with a V.35 interface; some support two wide area links per bridge.

This evaluation examines the bridges' installation, configuration, and administration procedures; adherence to ISO standards; performance; and features. An emphasis is placed on the performance of the bridges given the task of bridging two network segments running Novell NetWare 386 3.1 across a wide area link. The wide area link is established using a pair of General DataComm 552 Channel Service Unit/Data Service Units (CSU/DSUs). In all cases, the bridges use an external timing source to synchronize data transfer over each wide area connection.

Filtering/Forwarding

Remote Ethernet bridge manufacturers often refer to compatibility and performance specifications in their product descriptions. Product specifications characterize performance in terms of Ethernet frame (or packet) filtering and forwarding rates. Published filtering and forwarding rates do not always accurately characterize actual bridge performance because they do not account for internal delays or latency.

Latency refers to the time it takes the bridge to examine and retransmit a packet. A bridge may be capable of forwarding packets at the theoretical Ethernet maximum of 14,881 packets per second (pps). Assuming the bridge has room to store incoming packets (frame buffer) while other packets are queued for transmission, only a small initial delay occurs as the first small group of packets is readied for transmission. After this initial delay, packets stream at the specified rate. Delays are typically less than a few milliseconds, but on some networks (Novell NetWare 386 in particular), all data packets sent by applications between server and workstation require acknowledgment frames and latencies are cumulative. The remote communications link can also constrain performance.

NSTL's test network configuration consisted of two 802.3 network segments with a high-speed, synchronous, full-duplex connection. This connection simulates T1 high-speed digital communication, which has a well-defined throughput limitation of 1.536M bps. If the incoming packet transmission rate exceeds the bridge's throughput rate to the WAN link, the incoming packets are stored in the bridge's frame buffer until transmission. A prolonged disparity in incoming and outgoing traffic rates can cause lost or dropped packets when the bridge's buffer fills and overflows. A deep frame buffer (i.e., higher memory allocation) helps a bridge accommodate periods of intense packet forwarding.

Bridge Applications

Traffic Isolation

When traffic on a network becomes intense and performance degrades substantially, businesses often divide or "segment" the network using a local bridge, isolating traffic on two local segments. Ethernet packets with source and destination addresses on the same segment do not appear on the other segment. Connecting two remote LAN segments achieves a similar effect, except for the introduction of a remote communications link. The network administrator is responsible for ensuring cost-effective use of potentially expensive communications facilities (such as a full or fractional T1 link) with limited bandwidths.

Bandwidth Enhancement

The traffic isolation obtained with Ethernet bridges can improve the available bandwidth in a system of multiple connected Ethernet segments. Multiple LAN segments can be connected linearly with local bridges, thus isolating traffic to its local segment. If three or more segments are connected in this way, traffic between the two endmost segments can lead to substantial congestion on the middle segments. A backbone network topology can be created where a single network segment serves as a "backbone" or "tree trunk," from which additional segments, connected via bridge devices, branch. Since the backbone is devoted to handling nonlocal traffic, it will tend to be less congested than each local segment. The expense and capabilities of different types of remote communications links have an impact on this type of extended network organization when LAN segments are connected via remote Ethernet bridges.

For example, a customer service organization within a large corporation might discover a need to connect several physically separate offices within a single metropolitan area, as well as a need to connect those offices to similar offices in another city. By examining information needs and data throughput requirements among these offices, it might be found that several types of remote links, based on a heirarchy of bandwidths, will allow all the offices to be connected in a cost-effective way. Different grades of Digital Dataphone Service (DDS) up to 56K bps may be adequate to handle nonlocal traffic volumes between local offices in one city, while one of the local offices, acting as a backbone, may require a more expensive digital leased line, such as T1 (1.536M bps of effective throughput), and thus effectively handle all nonlocal traffic from one set of local offices to another set in another city.

Fault Tolerance

In addition to providing electrical isolation between segments, multiple bridges can provide routing fault tolerance with redundant network connections. If an active bridge fails, a backup or standby bridge ensures communication between segments. Some remote bridge devices have this sort of fault tolerance built in. If the remote bridges can be configured with more than one WAN port per bridge, the second wide area port can be used as a redundant wide area link to guard against communication interruptions should one link fail. In addition to enhancing communications reliability, some multiple WAN port bridges make effective use of the additional bandwidth with a second link. Remote bridges with fallback dial-up automatically dial into a leased digital communications facility via the second WAN port when the bridge's primary link fails. Fallback dial-up establishes connection with little or no network interruption.

Load Balancing

A well-planned traffic isolation scheme uses one or more bridges to balance the traffic load on interconnected network segments. In some cases, priorities can be assigned to individual bridge channels with priority levels relative to other bridges. Load balancing features give the administrator some control over the path of traffic across the extended network. Remote Ethernet bridges with multiple active WAN ports can sometimes share and balance loads among their WAN channels to effectively enhance overall WAN performance.

Enhanced Security

Bridges are no substitute for well-planned network security at the operating system and application levels, but they do offer additional security. Bridges placed on- or off-line can completely control access to critical resources, and most bridges offer complete control over frame filtering and forwarding according to specific network addresses. A particular node can be denied access to a specific segment given proper bridge administration, but in a large extended network, such administration may not be feasible. Many bridges can restrict forwarding of frames based on higher level protocol information contained within the frame, and some support filtering based on wild card symbols specified via the bridge administration facility. All the test bridges can learn the network addresses of devices on connected segments and thereby build address tables with several hundred or several thousand entries. An address database of this size would require extensive manual editing to achieve node-specific security arrangements. More generalized security measures such as protocol filtering are easily implemented.

Mixed Media Connection

Most remote Ethernet bridges can connect network segments with different physical media. All the test bridges can connect thin Ethernet to thick Ethernet cabling. Because remote bridges operate at the Data Link layer of the ISO model, they are inherently protocol and operating system independent. Physically separated LANs employing different media (thin Ethernet, thick Ethernet, twisted-pair Ethernet, or fiber optic Ethernet) can be connected with remote Ethernet bridges, assuming that appropriate local interfaces and connections can be made. In some cases, each bridge can be configured with the necessary local network interface; otherwise, an external transceiver may be required.

Physical Media Extension

Bridges are similar to repeaters in that they receive and retransmit network frames, providing a degree of electrical conditioning to the transmitted frame's signal. Electrical conditioning and the topology enhancements described above can greatly extend the media lengths allowed with a single network segment. With remote bridges, the distance separating local segments in a WAN is limited only by the remote communications link or links.

Summary

An important consideration when selecting a remote Ethernet bridge is the type and cost of the remote connection the bridge will use. While throughput and latency performance indicators should be given serious attention, keep in mind that features are as important as raw performance in specific environments. Performance is especially important with regard to expensive digital transmission services, and some products make more effective use of the remote link's available bandwidth.

The Magnalink does an outstanding job of cramming data into the link using a single WAN port. A pair of Magnalink bridges costs more than twice as much as a pair of Hewlett-Packard bridges, but 4-to-1 data compression and support for the spanning-tree algorithm give the Magnalink the best overall performance. With expected support for SNMP management standards, free customer support, flexible filtering capabilities, and fallback dial-up support, the Magnalink offers a powerful combination of performance and features.

The inexpensive Hewlett-Packard provides excellent overall performance, especially for a bridge with a single WAN port. It does not support fallback dial-up, but it does support the spanning-tree algorithm for redundant links. Its limited 512-entry address table is insufficient for networks with very large numbers of nodes. The Hewlett-Packard provides "plug-and-play" remote Ethernet bridging; its excellent documentation makes installation a snap. SNMP support and an optional Windows-based management package enhance usability.

The Retix offers two WAN links operating at T1 speeds with effective load balancing and load sharing. The Retix supports more management standards than any of the other bridges. The Retix excels at overall throughput using two high-speed digital links under a wide range of traffic conditions. Redundant links with very good throughput provide excellent fault tolerance. The Retix does not support fallback dial-up.

The solid 3Com product supports the spanning-tree algorithm and SNMP management standards, and it can be optionally configured as a router. Good performance combines with a robust set of features at a competitive price.

The very expensive Vitalink comes with high levels of supplier support and commitment. The Vitalink bridge supports up to eight WAN ports per bridge with fallback dial-up at rates to 384K bps. Very good features, including flexible filtering, make the Vitalink a very good product. Overall throughput for a double WAN port bridge using two T1 links is the highest tested; intermittent channel shutdown reduces throughput accordingly.

The RAD comes configured with two WAN ports and can be expanded to four ports per bridge. It stores up to 25,000 source address table entries on its internal hard disk. Vitalink's excellent bridge management software must be purchased separately. Very flexible filtering abilities, spanning-tree support, and source routing transparency make this PC-based remote bridge viable for many internetworking needs. At $8,900 with the RIM software, the RAD is somewhat pricey compared to bridges with better performance. A more powerful PC platform would undoubtedly improve performance (tested on an AT-class PC platform operating in nonturbo mode).

Product Evaluations

Hewlett-Packard HP 28674A

Product Summary

- Excellent performance under all traffic conditions
- Very good throughput (single WAN port)
- Very small internal latencies
- Linear relationship between latency and packet size
- Excellent documentation; easy installation
- Console command line management software
- Automatic learning of source addresses
- Up to 512 source address table entries
- Supports spanning-tree algorithm, SNMP
- Small footprint
- Rack-mount hardware included
- Maximum 2.048M bps WAN link speed
- External timing source required for WAN link synchronization
- Windows-based management software available

Magnalink 3000 Series HCB

Product Summary

- Excellent performance under all traffic conditions
- Excellent throughput; rivals bridges with two WAN links
- Consistently small internal latencies
- Somewhat nonlinear relationship between latency and packet size
- Easy installation
- Console command line management software
- Automatic learning of source addresses
- Maximum 8,192 source address table entries
- Supports spanning-tree algorithm, SNMP (third-quarter 1991)
- Free customer support
- Maximum 2.048M bps WAN link speed
- Rackmount integral with unit case
- Fallback dial-up to 56K bps

RAD REB-20

Product Summary

- Good performance with moderate traffic; significant degradation with constant heavy traffic
- Linear relationship between latency and packet size
- Relatively high internal latencies
- Good documentation; easy installation
- Optional menu-driven management software; powerful and easy to use
- Optional NetGraph software provides graphic representation of network topology and status
- Management software easy to install; can access bridges from workstation
- Automatic learning of source addresses
- Maximum 25,000 source address table entries
- Free customer support
- Maximum 2.048M bps WAN link speed
- Fallback dial-up to 4800 bps

Retix High Performance 4880

Product Summary

- Consistently excellent performance for all traffic conditions
- Excellent throughput; effectively uses two WAN links at T1 speeds
- Automatic learning of source addresses
- Maximum 16,000 source address table entries
- Supports spanning-tree algorithm, SNMP, MAP, TOP
- Effective menu-driven bridge management software
- Robust configuration and administration facilities
- Easy installation
- Effective load balancing and sharing on two WAN links
- Maximum 2.048M bps WAN link speed
- Optional Windows-based management software
- Free customer support

3Com NETBuilder IB/3000

Product Summary

- Very good performance with moderate traffic
- Performance degrades with constant heavy traffic
- Good throughput; low internal latencies
- Linear relationship between latency and packet size
- Comprehensive documentation; easy installation
- Automatic learning of source addresses
- Maximum 8,196 source address table entries
- Load sharing on two WAN ports
- Console command line management software
- Supports spanning-tree algorithm, SNMP
- Optional router software
- Maximum 7M bps WAN link speed
- External or internal WAN link timing
- Fallback dial-up

Vitalink TransLAN 335

Product Summary

- Good performance under all traffic conditions
- Excellent throughput with two WAN ports active; shuts down one port under heavy loading
- Good throughput with one active WAN port
- Relatively high internal latencies
- Linear relationship between packet size and latency
- Very good documentation; easy installation
- Robust configuration and administration environment
- Console command line management software with rudimentary user interface
- Supports spanning-tree protocol
- Load balancing and sharing among WAN ports
- Maximum 2.048M bps WAN link speed
- Fallback dial-up to 384K bps
- Up to eight WAN ports per bridge

Product Recommendations

Magnalink 3000 Series HCB

The Magnalink earns NSTL's top recommendation for its strong performance and effective use of available remote link bandwidth. Four-to-one data compression at T1 speeds, support for the spanning-tree algorithm, expected support for SNMP management standards, free customer support, flexible filtering capabilities, and fallback dial-up support make the Magnalink a powerful, high-performance product. The Magnalink with one wide area link provides throughput equal to what many other bridges provide with two WAN links.

Hewlett-Packard HP 28674A

The Hewlett-Packard offers excellent overall performance at a great price and an extremely easy installation procedure. The Hewlett-Packard supports the spanning-tree algorithm for redundant links between LAN segments using multiple bridges. It does not support fallback dial-up, and its 512-entry address table may be too limiting for some extended networks. SNMP support and optional Windows-based management software enhance usability. Hewlett-Packard has a reputation for high-quality products, and the Hewlett-Packard remote bridge is no exception.

Retix High Performance 4880

For overall throughput using two high-speed digital links under a wide range of traffic conditions, the Retix may be the best choice. The Retix effectively uses two WAN links operating at T1 speeds with load balancing and load sharing, and it supports more management standards than any other bridge tested. Redundant links, both with very good throughput capabilities, provide excellent fault tolerance. Fallback dial-up is not supported by this spanning-tree bridge. Retix has a solid background in internetworking products.

Rating Summaries

Figure 1.
Hewlett-Packard HP 28674A Ratings

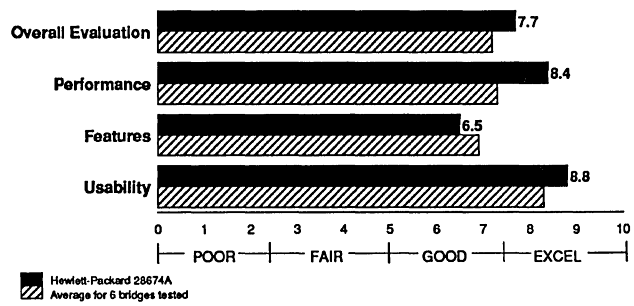

Figure 2.
Magnalink 3000 Series HCB Ratings

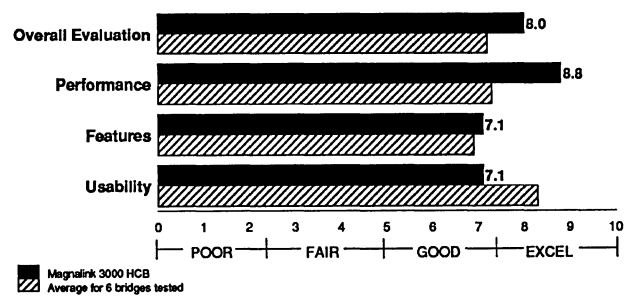

Figure 3.
RAD REB-20 Ratings

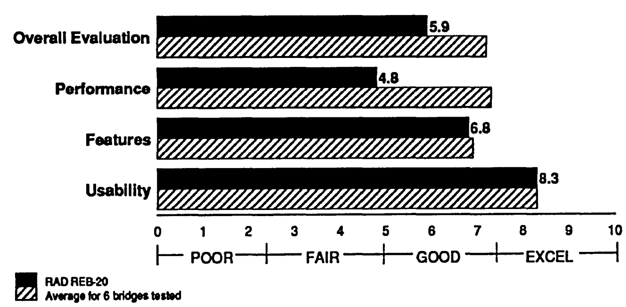

Figure 4.
Retix High Performance 4880 Ratings

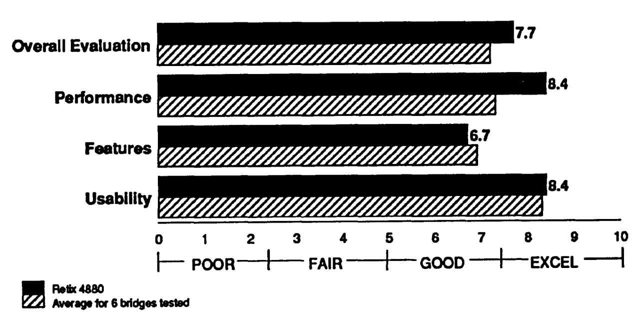

Figure 5.
3Com NETBuilder IB/3000 Ratings

Figure 6.
Vitalink TransLAN 335 Ratings

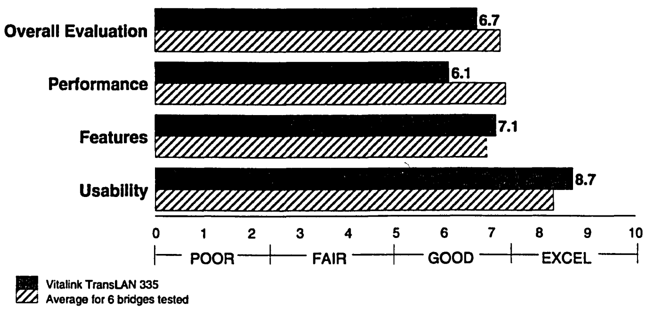

Overall Evaluation

The bridges are uniformly easy to install and operate, and features are generally very good. Performance and features generally distinguish the test products. For heavy network traffic and potentially complex network management scenarios, good performance and bridge management capabilities are a must. Products with strong performance characteristics and robust configuration and management features can provide the network manager with tools to enhance the organization and performance of an extended network, but their effectiveness depends upon the way they are integrated with overall network management strategies.

Magnalink's excellent performance, very good features, and reasonably good usability result in a winning combination. Its capability to stuff tremendous amounts of data into a single WAN link through data compression can reduce long-distance communications costs (e.g., a fraction of a T1 line may meet the information needs of an entire organization). Fallback dial-up capability that currently supports speeds to 56K bps (with support for T1 speeds planned) can provide a backup link within seconds of a primary link failure.

Hewlett-Packard and Retix offer good features and excellent performance and usability. The Hewlett-Packard has a single wide area port per bridge, and Retix has two per bridge. Retix uses a significant portion of the bandwidth provided by two T1 links and does so flawlessly. Hewlett-Packard is not capable of Retix' throughput by nature of its single WAN link, but it makes impressive use of the available bandwidth with a single WAN link and virtually fills the T1 bandwidth. Even with sustained, intense background traffic, the Hewlett-Packard is not limited by its throughput.

The 3Com combines very good performance, very good features, and excellent usability, but a bridge with two WAN ports, load balancing, and load sharing should provide better throughput. Throughput curves for two T1 links make it difficult to believe that the higher remote link speeds supported by the 3Com are fully utilized. With larger packets the effective data throughput (in bytes per second) might increase markedly, but within the context of NSTL's tests, the 3Com should be capable of pushing more packets through the links. This theory is borne out by the 3Com's slightly depressed scores with flood traffic. Performance is decent, and the bridge can function as a router with the appropriate software. With improved remote link throughput, the 3Com could be a real winner.

The Vitalink offers good features, good performance, and excellent usability. While exhibiting the highest throughput of all the bridges tested, the Vitalink intermittently shuts down one of its wide area channels during periods of sustained intense packet forwarding. The channel shutdown, combined with slightly higher latencies than most of the other bridges, may be the cause of some of Vitalink's slower test completion times and its sensitivity to flood traffic. This channel shutdown effect is intermittent, occurs only under sustained intense loading, and disappears when loading eases. Vitalink expects to correct the problem in the near future.

The RAD provides good features, excellent usability, and fair performance. RAD's optional Remote Internetwork Management software makes bridge administration a breeze. Without the software, the RAD cannot be administered during operation. Like the other double-WAN port bridges, the RAD should provide greater throughput; in fact, its throughput significantly slows in the presence of constant intense traffic (flood traffic). Load sharing and balancing do not seem to help in NSTL's tests. The RAD's performance may be constrained by the performance limitations of its AT-class PC platform.

Methodology

The Overall Evaluation is a weighted average of scores for the individual criteria.

Overall Evaluation Score = (5 x Performance Score) + (4 x Features Score) + (Usability Score) ÷ 10.

Performance

Small delays (known as latency) in the forwarding of 802.3 frames from one network play a large role in determining performance. Expressed in the form of latency curves, the relative latencies of the test products serve as a reasonably good indicator of expected bridge performance. Latency increases with increases in packet size, and in most cases the relationship is linear.

The Magnalink uses a 68030 processor with a proprietary system and bus architecture. Using a single wide area link, it accomplishes throughput expected only from a bridge with two WAN links. The performance is impressive and translates into lower remote communications costs. With Magnalink bridges, an organization's data needs may be met using a single T1 line rather than two lines, or using a smaller fraction of a T1 line. Magnalink accomplishes high throughput using 4-to-1 data compression at T1 speeds, and in some instances it achieves almost 5-to-1 compression (based on statistics gathered by the bridge). Low overall latencies contribute to the Magnalink's excellent performance.

The Retix and Hewlett-Packard also offer excellent performance. The Retix bridges exhibit tremendous throughput, fully utilizing two WAN ports per bridge. The

Hewlett-Packard achieves very good throughput, especially for a bridge with a single WAN port. Throughput is not a limiting factor for the Hewlett-Packard bridges, and their excellent performance reflects the capability to fill the remote communications channel. The Hewlett-Packard and Retix have low overall latencies.

The 3Com provides very good performance with two WAN ports. Throughput is good and latencies are low overall, but throughput levels are not as high as might be expected with two WAN ports, load sharing, and load balancing.

The Vitalink achieves the highest level of throughput, but intermittent shutdown of one wide area link reduces throughput to the level of a remote bridge with one WAN port. Vitalink expects to correct the problem in the near future. The Vitalink bridge pairs exhibit slightly higher latencies than several other bridges.

The RAD's performance is not quite as good as the Vitalink's, and its latencies are the largest overall. The RAD experiences difficulty with flood traffic, and the constant traffic intensity pushes it to the limits of its throughput capability. Some of the RAD's performance limitations may be the result of its relatively slow 80286 platform.

Methodology

For each iteration of each benchmark, individual performance scores are weighted and scaled relative to the best time for that component to produce a System Performance Index. Overall Performance scores are weighted averages of the System Performance Indexes for all benchmarks.

Performance Score = (Latency Score) + (Lotus 1-2-3 Score) + (Microsoft C 5.0 Score) + (Foxbase+/LAN Score) + (Xcopy to Server Score) + (Xcopy from Server Score) + (Throughput Score) ÷ 7.

Features

All the remote Ethernet bridge products provide robust features. The 3Com stands apart with support for over 8,000 source address table entries, remote communications up to 7M bps (not tested), and a rich set of administration and configuration features. 3Com builds in all sorts of forwarding and filtering features, good system security, and support for the spanning-tree protocol and Simple Network Management Protocol (SNMP). Source addresses are automatically learned and aged. As with all good MAC-layer bridges, the 3Com's operation is transparent to various operating systems and higher layer protocols while providing a wide range of protocol- and operating system-specific filtering.

The Magnalink and Vitalink are MAC-layer bridges with operating system and protocol transparency. Vitalink excels in bridge statistics reporting but does not support SNMP. Magnalink plans to provide SNMP support in the third quarter of 1991. Both bridges support the IEEE 802.1 spanning-tree protocol. Both Vitalink and Magnalink support remote link speeds up to 2.048M bps, provide fallback dial-up on a second remote line, and support over 8,000 source address table entries.

The RAD and Retix features are closely matched. Retix supports 16,000 source address table entries; RAD supports 25,000 by virtue of its capability to use hard disk storage. RAD supports a wider range of operating system- and protocol-specific filtering. Both bridges support the spanning-tree protocol, and Retix also supports SNMP, MAP (Manufacturing Automation Protocol), and TOP (Technical Office Protocol). Both bridges have two wide area ports per bridge, and the RAD supports fallback dial-up. RAD's REB-40 model (not tested) supports up to four WAN ports. Both Retix and RAD provide a full range of administration and configuration utilities. Like the other test bridges, the Retix supports leased lines, AT&T DDS, T1, and fractional T1; but unlike the others, it does not support Switched 56 (56K-byte dial-up).

The Hewlett-Packard supports up to 512 source address table entries, has a single wide area port per bridge, and supports SNMP and the spanning-tree protocol. It does not provide fallback dial-up, but it supports remote link speeds to 2.048M bps. Good system security, a full range of protocol and operating system security, and secure-mode filtering (pass only specified type) are included.

Methodology

NSTL compares each remote Ethernet bridge against a master list of features and specifications. The Features score is a weighted average of scores for the individual weighted features. Features and their methodology weights are listed in Table 3.

Usability

The remote Ethernet bridges are easy to install and tend to work well with their factory default settings. Cabling requirements are slightly more complex than for local bridges, depending on the remote communications link and the bridge's WAN interface. Setting the appropriate serial communications parameters is sometimes more complicated for bridges with more than one WAN port per bridge, especially when load sharing and load balancing settings needed to be adjusted.

The Hewlett-Packard is very simple to use and features an automatically configuring WAN port. The bridge has a small footprint, is very lightweight, and comes with rack-mount hardware. LED indicators on the back provide reasonably good diagnostic and operational information. User manuals are well organized, comprehensive, and professionally produced. Bridge management uses console commands with a rudimentary but adequate user interface. Windows-based management software is available as an option.

Retix and RAD supply well-organized management software with somewhat more sophisticated menu-driven user interfaces than the others. RAD's Remote Internetwork Management (RIM) software is very effective and easy to use, and NetGraph displays graphic representations of network topology and status; both must be purchased separately. (The basic RS-232 version of RIM costs $1,950.) RIM is almost a requirement for the RAD because without it the bridge cannot be administered during operation. The Retix user manual is very comprehensive, but it lacks an index and could benefit from better organization and a hardware/software checklist. RAD's user manual lacks an index and falls short in comprehensiveness. Good context-sensitive Help in the RIM software partially makes up for this shortcoming. Both bridges are easy to install, but the RAD requires additional space if both the monitor and keyboard remain attached to the PC platform. The Retix' LCD front-panel display provides detailed operational and diagnostic information. Retix also offers Windows-based management software.

The Vitalink comes with excellent manuals, installs very easily according to step-by-step instructions, and despite its rudimentary user interface, the management software is well organized and powerful. The large-footprint

Vitalink has a low profile and does not require much vertical space. LED indicators on the back of the bridge provide status and diagnostic information.

The 3Com comes with an excellent Getting Started manual. The draft version of the user guide is bulky, lacks an index, and is difficult to use although extremely comprehensive. The manuals provide no hardware/software checklist for facilitating installation. Overall, the 3Com is not difficult to use considering the depth and versatility of its features (router software operates on the same platform). Its management software interface is simple, but unsophisticated.

Magnalink ranks slightly lower than the other bridges because its manual lacks comprehensiveness, organization, and an index. A better page layout and type style and a hardware/software checklist would improve the documentation considerably. Despite the quality of its manuals, the Magnalink is easy to install and use. The bridge management software uses a rudimentary command line interface and provides the user with the tools for adequate bridge configuration and management. Front-panel LED indicators display operational status and diagnostics. The front panel removes easily, revealing additional LEDs and the main circuit boards. The bridge chassis integrates rack-mount hardware.

Methodology

The Usability score is a weighted average of scores for the individual criteria.

Usability Score = (3 x Physical Installation Score) + (2 x Manual Organization Score) + (2 x Manual Clarity Score) + (3 x Manual Comprehensiveness Score) + (2 x Console/Administration Hookup Score) + (3 x Console/Administration Use Score) ÷ 15.

Performance Results

Manufacturers of remote Ethernet bridge products characterize performance in terms of the maximum packet filtering and forwarding rates and types of remote connections supported. Measured with dedicated test equipment in a controlled Ethernet environment, these "raw" filtering and forwarding rates do not completely describe the behavior of remote Ethernet bridges. Many factors contribute to the "real-life" performance that can be expected from a pair of remote Ethernet bridges. For bridges that support high-speed serial communications (e.g., over T1 lines), frame buffering and transfer between the LAN and the wide area link (or links) are critical to performance.

High-speed serial communications lines can be expensive, and organizations using the entire T1 bandwidth need a bridge capable of filling the bandwidth. Remote Ethernet bridges adequate for communication over DDS (maximum 56K bps synchronous digital transmission) may not be capable of keeping pace with one or more 1.544M bps (1.536M bps plus framing bits) lines.

As with local Ethernet bridges, packet size, queuing methods, interpacket spacing, bandwidth utilization, and internal bridge latency play important roles in remote bridge performance. Since remote Ethernet bridges must operate in pairs connected via WAN ports, many other factors contribute to overall system performance. Latencies associated with moving packets from one LAN to another via a synchronous, full-duplex digital carrier (e.g., T1 line) depend on the speed of the line as well as the bridges' internal processing time at either end. Round-trip delays may be inherited from the carrier's digital network.

Test Configuration

NSTL tested remote Ethernet bridges connecting two 802.3 networks. Each network used thin Ethernet cable and 10BASE2 Ethernet cards; the bridges were connected by means of a simulated T1 line. Each bridge in the pair connected to one network and to a pair of General DataComm 552 CSU/DSUs. Twisted-pair cables with appropriate crossovers connected the CSU/DSUs, each configured to provide communication at T1 speeds. One CSU/DSU (the Master CSU/DSU) provided timing for the simulated T1 line. The remote bridges were configured to look to an external source (i.e., the Master CSU/DSU) for timing over the remote communications line. Bridges with two wide area ports were connected with two pairs of CSU/DSUs.

Each network consisted of one server running Novell NetWare 386 3.1, four secondary workstations, and one primary workstation. On one LAN, a control workstation and monitor workstation ran Network General's Watchdog software; a Network General Sniffer was attached to the other LAN. All workstations accessed logical network drives. Control stations had no effect on test results; they monitored and recorded network traffic conditions during testing and provided global traffic statistics and frame distributions for the entire network and for individual workstations.

Secondary Workstations: Everex 386sx systems identically configured with 4M bytes of RAM, Novell NE2000 Ethernet adapters.

Primary Workstation: PC Craft 386/33 Desktop with 4M bytes of RAM, Western Digital 16-bit EtherCard Plus 16.

Control Station: AST Bravo/286 with 640K bytes of RAM, Western Digital 16-bit EtherCard Plus 16.

Network Server: PC Craft 386/33 Tower with 12M bytes of RAM, 80M-byte IDE drive, Western Digital 16-bit EtherCard Plus 16.

Monitor Workstation: AST Premium 386sx/16 with 4M bytes of RAM, 40M-byte hard disk drive, Racal InterLan NI5210 DataLink-Level Ethernet Controller, Network General Watchdog 1.00.

CSU/DSU: General DataComm 552 Data Service Unit.

Network Analyzer: Network General Sniffer, Compaq Portable 386/20, Network General's Ethernet Analyzer 3.0, and Ethernet Monitor Software.

Application Tests

Four application benchmarks measured bridge performance with typical Lotus 1-2-3 Release 3, Microsoft C, Foxbase, and Xcopy operations. As a baseline measurement, the No Traffic times were measured from the primary workstation with the secondary workstations idling. Background application traffic was generated using Lotus 1-2-3, a Foxbase+/LAN application suite, a Foxbase+/LAN traffic application (Spike Traffic), and a sustained intensity low-level traffic generator (Flood Traffic).

High-level application tests with and without background traffic ran symmetrically on the two networks. Applications ran simultaneously at the primary workstation on each segment with the appropriate traffic conditions established on each segment's four secondary workstations. To maintain balanced loading on both network segments, all primary and secondary workstations were logged into NetWare servers via the bridge pairs; therefore, all information transfer over the network passed through the bridges.

Internal Latency Tests

Remote Ethernet bridges are said to operate in pairs because each segment requires its own bridge. The remote bridges are then connected via their wide area network ports using long-distance communications lines. (In some cases, a single bridge with multiple WAN ports can be connected to several remote bridges.) Packets transferred from one segment to the other must cross two bridges and be processed at both ends. At either end of the wide area link, frames are delayed by processing required to place them on the local segment. Communications equipment may add some latency to the system, but the same equipment is used for all the bridge pairs, so relative comparisons among bridges are valid. A bridge's processing speed is determined by its processor type and speed, quantity and type of RAM, the number of address entries maintained, its method of searching for addresses, queuing of packets for transmission, packet spacing, and the efficiency of algorithms used to orient the bridge on the network.

Internal latency is the time required for the bridge to process a received Ethernet packet and forward or filter the packet based on the destination address. Upon receiving a packet, the bridge references its internal address database to determine whether the destination is on the same segment as its source. Packets that remain on the current segment are filtered; others are forwarded across the remote communications link to the other bridge where it is again processed and moved onto the other segment. A filtered packet essentially remains local and is received at the destination address. Bridge processing is minimal if the internal address database is fairly small and few address comparisons are required. Packet transmission and receipt on the local segment proceed independently of bridge activity, so bridges do not introduce any delay between stations on the same segment.

Forwarding packets is essentially more time consuming than filtering (analyze and retransmit versus analyze and do nothing), and other factors add to the delay. Packets queued for forwarding in bridge memory contribute to retransmission latency. The speed of the communications link (wide area link) can bear directly on transmission delays, depending on how each bridge handles queued packets.

NSTL treated the two remote Ethernet bridges, CSU/DSUs, and synchronous full-duplex connection as one system with complex internal interactions. By treating the system as a whole and holding certain factors constant (CSU/DSU speed and timing, LAN equipment, and benchmarks), NSTL effectively measured and compared latencies for each bridge pair and its remote link.

Internal latency results measured the time required to forward a large number of packets across the bridge. Two nodes, one on each segment, conducted an IPX conversation using custom send and receive programs optimized for throughput with no disk access or extraneous data processing to hinder transmission. The IPX send and receive programs provided address-specific packet transmission and dropped packet time-out checking to ensure the successful forwarding of all packets. Each conversation was repeated five times with five packet sizes from 100 bytes to 500 bytes, providing highly accurate results. In each conversation, 3,000 packets were forwarded. The times were large enough that errors in measurement were approximately 1% or less.

For comparison, all conversations were repeated without a bridge (Null Bridge in the charts); that is, the segments were joined into one network. Internal latency was measured without network traffic. The two primary workstations served as the sender and receiver stations.

Throughput Tests

NSTL's throughput test measured the bridge's capability to fill the available bandwidth or data channel. For testing, the CSU/DSUs connecting the bridge pairs provided a 1.536M bps (T1 speeds) effective bandwidth per remote link. Bridges with two wide area links used two remote links, for an available bandwidth of 3.072M bps. In addition to system latencies, frame buffer size, buffering method, data compression, and load sharing and balancing among WAN ports also affected sustainable throughput. Extreme traffic conditions often revealed throughput limits, as shown in the benchmark results.

Limitations on a bridge's effective forwarding rates can result in dropped packets and reduced throughput. Higher level protocols generally initiate retransmission of dropped packets. When packet loss is minimal, retransmission delays are negligible. If packet loss increases sufficiently, the cumulative delays can noticeably degrade network performance.

Packets were transmitted across the bridges to a secondary workstation on the other segment. By sending packets at a range of speeds (measured in packets per second, or pps), it was possible to determine the packet transmission rate at which each bridge pair can no longer forward packets as fast as it receives them. The curves in Figure 11 show throughput trends for the tested bridges. Packet transmission rates were measured on the source segment using a Network General Sniffer, and on the destination segment using a monitor station running Watchdog 1.0.

Background Traffic

NSTL generated high- and low-level background traffic to measure its effects on bridge performance. High-level applications such as Lotus 1-2-3 and Foxbase+/LAN rely on the operating system and its higher level protocols to send and receive data. Secondary workstations on each segment were logged into servers across the bridge to provide balanced network loading. Low-level traffic generated using custom programs created traffic conditions not typically found with high-level applications. NSTL's low-level Flood traffic transmitted a constant stream of minimum-size packets to a specific address and did not require acknowledgement frames via higher level protocols.

High-level background traffic was generated by Lotus 1-2-3 and Foxbase+/LAN. Simple Foxbase background traffic consisted of each secondary workstation performing a transaction no more than 30 times a minute. The transaction looks up indexed records in three files and locks, updates, and unlocks one record in each file. Spike traffic, also generated by Foxbase+/LAN, consisted of each secondary workstation repeatedly updating a specific record in a single database file, remaining dormant for a brief interval, and then repeating the cycle with a different number of updates and a different delay time. Lotus traffic executed a macro that combines and saves files metered to be

Figure 7.
Lotus 1-2-3

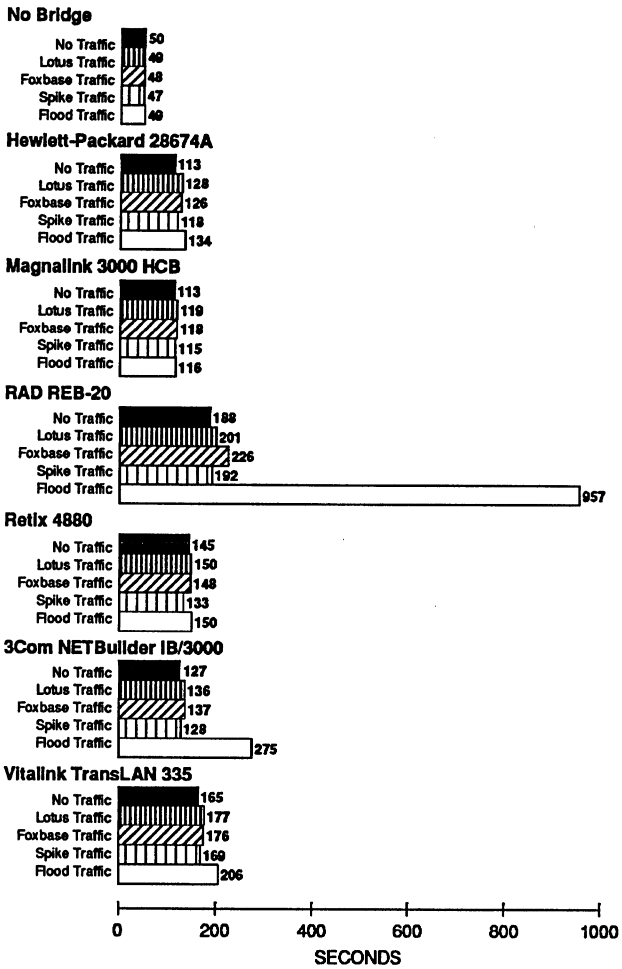

Figure 8.
Microsoft C 5.0

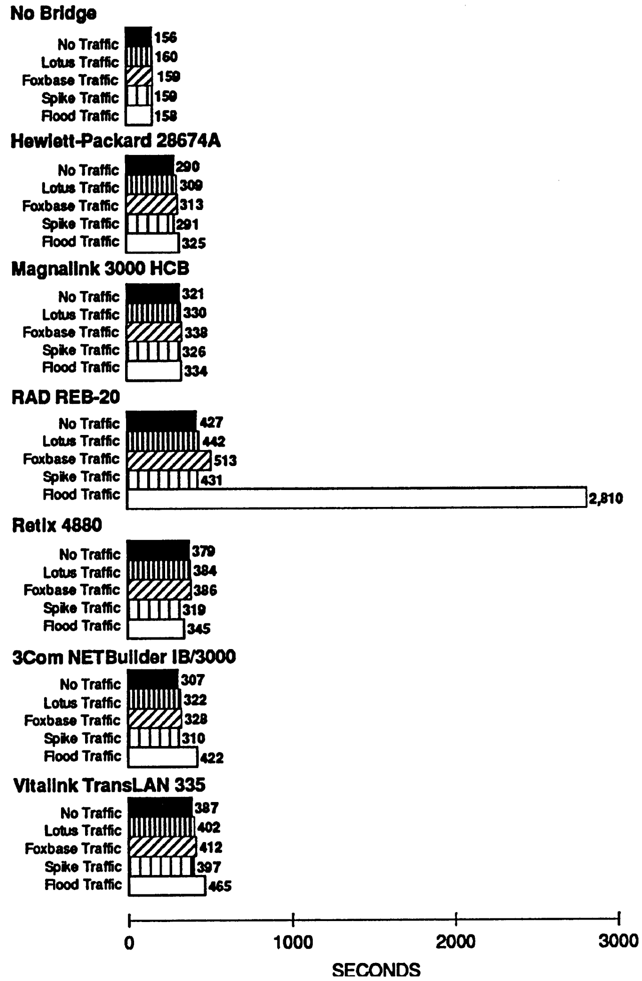

active no more than every 20 seconds on each station. As configured in the test bed, Lotus traffic provided a level of background network utilization ranging from 1% to 7%. Foxbase traffic network utilization ranged from 3% to 8%. Spike traffic is similar to Foxbase traffic, but network utilization ranged from less than 1% to about 8%. Half of the Spike traffic frames fell within the 129- to 256-byte range; the other 50% were 60-byte frames. In Foxbase traffic, about 55% fell in the 61-to 128-byte range, and the others fell into the 60-byte, 129-to 128-byte, and 513- to 1,024-byte ranges.

Low-level Flood traffic was optimized for throughput and continuously sent minimum-sized Ethernet packets to a specific network address at a controlled rate. To ensure symmetric loading of the segments, two secondary workstations on each segment sent packets to two idling secondary workstations on the opposite segment. Each of the four secondary workstations generated packets at 600 pps, equivalent to about 3% network utilization.

Lotus 1-2-3

A Lotus macro invoked on both primary workstations loads and saves a 3M-byte spreadsheet (to and from a server), moving a substantial amount of data across the network. This test provides a well-balanced data transmission scenario between workstation and server. The test is run without background traffic, then with four types of background traffic running concurrently on eight secondary workstations, four on each network segment.

Analysis

The results group the bridges into two categories that roughly coincide with observed latency trends. Performance trends among the bridges are consistent with high- and low-level traffic. Flood traffic accentuates these trends and indicates the effectiveness of the bridges' frame buffering and processing. Flood traffic affects the Magnalink least; the RAD exhibits significant performance degradation. The 3Com, Hewlett-Packard, Retix, and Vitalink slow some with Flood traffic.

Microsoft C 5.0

A large ensemble of 25 XLISP 2.0 C source code files (10,928 lines totaling 245,222 bytes) is compiled and linked at each primary workstation, forming an executable file of 147,228 bytes. The total number of bytes downloaded from each remote server is far greater than the amount of data traveling from workstation to server. The test is run without background traffic, then with four types

Figure 9a.
Xcopy to Server

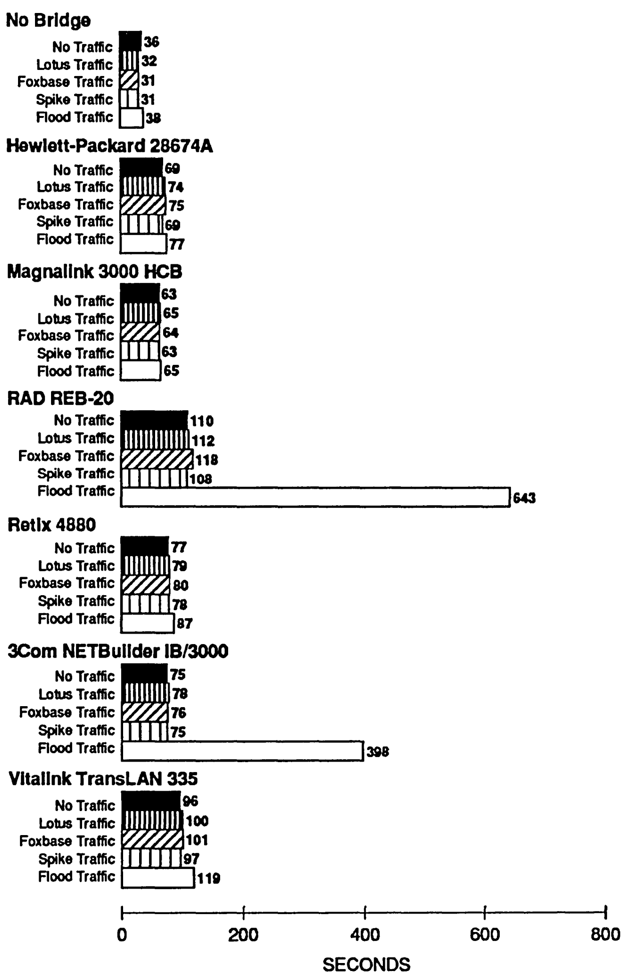

Figure 9b.
Xcopy from Server

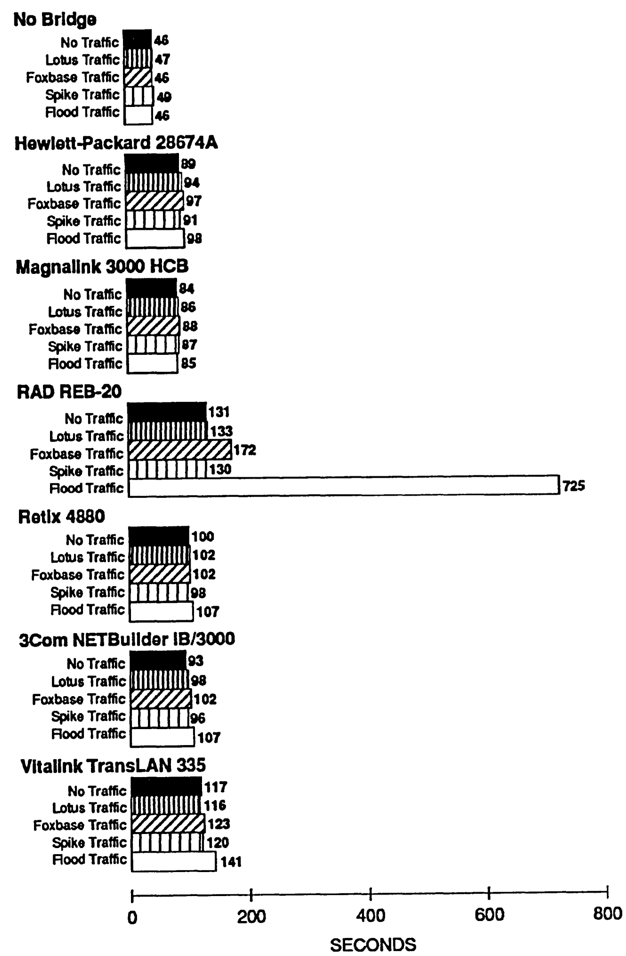

of background traffic running concurrently on eight secondary workstations, four on each network segment.

Analysis

Performance results coincide with latency trends. Times are longer than in the Lotus test because more files are downloaded and because the time-consuming compiling and linking cannot proceed until the files are downloaded. Bridges that transfer the files to the workstation quickly have a recurring advantage.

The Microsoft C test transfers a larger percentage of small packets than the Lotus test, and the Hewlett-Packard bridge exhibits noticeably shorter latencies for smaller packets. Throughput apparently is not a limiting factor because the Magnalink and Retix have greater throughput capacities than the Hewlett-Packard.

The RAD's large latencies explain its slow times in the no traffic and high-level traffic scenarios, but its Flood traffic time reveals throughput limitations that can lead to dropped packets and subsequent retransmissions.

Xcopy

A 5M-byte directory tree (130 files in 13 directories) is copied to each remote server from its primary workstation using the DOS Xcopy command with the /s parameter. The test is run without background traffic, then with four types of background traffic running concurrently on eight secondary workstations, four on each network segment. The entire test suite is repeated, copying the directory tree from the server back to its primary workstation.

Analysis

The Null Bridge results show that the Xcopy proceeds systematically about 50% slower from the server than to the server. The trend holds true with all types of traffic, except for the 3Com with Flood traffic, which reverses the trend.

Without traffic and with high-level traffic, the Hewlett-Packard, Magnalink, 3Com, and Retix achieve shorter completion times than the other bridges. Under similar traffic conditions, the Vitalink and RAD are slowed by larger latencies, and with Flood traffic, the 3Com, RAD, and Vitalink are slower.

Table 1. Foxbase+/LAN Tests

No Traffic (times in seconds)

	No Bridge	Hewlett-Packard HP 28674A	Magnalink 3000 HCB	RAD REB-20	Retix 4880	3Com NETBuilder IB/3000	Vitalink TransLAN 335
Test 1	61.0	69.8	69.3	77.0	85.3	70.8	75.5
Test 2	11.0	24.5	23.8	36.8	31.3	25.5	32.5
Test 3	16.0	43.0	37.8	67.5	53.5	45.5	58.3
Test 4	22.0	72.3	63.3	112.8	84.5	75.8	98.5
Test 5	8.0	22.5	21.5	33.5	26.3	23.8	30.0
Test 6	19.0	50.3	45.0	77.3	60.5	52.8	67.5
Test 7	21.0	27.0	28.5	33.3	34.8	28.5	32.5
Test 8	4.0	9.0	9.0	13.8	11.0	9.5	12.3
Test 9	31.0	60.3	64.5	87.0	77.5	66.8	81.5
Test 10	45.0	95.0	90.0	137.3	118.5	99.3	123.3
Test 11	40.5	94.0	101.0	145.3	118.5	102.0	135.3
Test 12	5.0	7.0	7.0	8.5	9.5	7.5	8.8
Test 13	11.0	23.8	22.0	35.8	30.5	25.3	31.5
Test 14	1.0	2.0	2.3	3.0	2.8	2.3	2.8
Test 15	1.5	2.8	2.5	3.5	3.3	2.8	3.0
Test 16	19.0	27.3	27.8	34.5	35.8	28.3	32.8

Lotus Traffic (times in seconds)

	No Bridge	Hewlett-Packard HP 28674A	Magnalink 3000 HCB	RAD REB-20	Retix 4880	3Com NETBuilder IB/3000	Vitalink TransLAN 335
Test 1	61.5	70.5	69.3	78.3	85.0	71.3	75.3
Test 2	12.0	26.0	24.0	36.8	31.0	27.0	33.5
Test 3	16.5	46.8	40.3	69.5	53.5	48.3	61.3
Test 4	23.0	78.5	66.3	117.0	85.8	82.3	102.8
Test 5	8.0	25.0	21.5	35.0	25.8	25.0	31.0
Test 6	18.5	55.3	46.8	79.8	61.3	56.3	70.5
Test 7	21.0	28.3	29.0	35.0	34.8	29.5	33.5
Test 8	4.0	9.5	9.0	14.3	11.3	10.0	12.0
Test 9	31.5	64.5	66.0	94.8	78.0	71.3	86.0
Test 10	46.0	101.5	92.8	143.3	121.0	105.3	128.0
Test 11	42.0	103.8	103.5	163.5	121.8	111.0	143.3
Test 12	4.5	7.3	7.0	9.0	9.5	7.3	8.5
Test 13	11.5	25.8	23.3	36.5	30.3	25.8	33.0
Test 14	1.0	2.0	2.3	3.8	2.8	2.8	3.0
Test 15	1.0	2.8	2.8	3.5	3.3	2.5	3.3
Test 16	19.0	28.0	27.5	37.3	35.0	29.3	33.8

Foxbase+/LAN

Foxbase runs a series of transactions against a banking database:

1. Write Indexes to Three Files
2. Create Indexes on Four Files
3. Copy Selected Records to Temporary Files
4. Delete Selected Records
5. Append Files
6. Pack Database (removing deleted records and reindexing)
7. Process 1,000-Transaction Batch
8. Create Indexes on History File
9. Process 1,000-Transaction Batch
10. Append Records to History File and Reindex

Table 1. Foxbase+/LAN Texts (Continued)

Foxbase Traffic (times in seconds)

	No Bridge	Hewlett-Packard HP 28674A	Magnalink 3000 HCB	RAD REB-20	Retix 4880	3Com NETBuilder IB/3000	Vitalink TransLAN 335
Test 1	61.0	69.8	69.5	79.3	85.8	71.3	76.0
Test 2	12.0	27.0	25.0	43.5	31.0	27.8	33.8
Test 3	16.0	47.0	40.5	78.0	53.0	49.8	61.8
Test 4	24.0	82.3	65.5	126.5	85.0	85.3	103.3
Test 5	7.5	24.8	22.5	38.8	25.8	25.8	31.3
Test 6	19.5	55.5	47.3	89.5	60.8	58.5	71.8
Test 7	21.0	28.0	29.0	37.3	34.5	29.8	33.3
Test 8	3.5	10.0	9.0	15.3	11.0	10.0	13.0
Test 9	32.5	63.5	67.0	101.5	77.3	70.8	86.0
Test 10	45.0	102.0	93.0	157.8	120.5	106.5	129.5
Test 11	41.5	101.5	104.5	171.8	117.5	110.0	144.5
Test 12	4.5	7.3	7.0	10.0	9.8	8.0	8.8
Test 13	11.0	25.3	23.5	40.5	29.5	26.8	33.3
Test 14	2.0	2.5	2.3	3.8	2.5	2.3	3.0
Test 15	1.0	2.5	2.8	4.0	3.3	2.8	3.0
Test 16	19.0	28.3	28.5	39.8	35.0	29.5	34.8

Spike Traffic (times in seconds)

	No Bridge	Hewlett-Packard HP 28674A	Magnalink 3000 HCB	RAD REB-20	Retix 4880	3Com NETBuilder IB/3000	Vitalink TransLAN 335
Test 1	71.5	71.0	70.0	69.0	77.8	61.0	75.3
Test 2	27.0	26.3	24.8	24.8	36.0	12.0	33.3
Test 3	48.0	46.8	44.0	39.5	67.0	16.5	60.3
Test 4	79.8	79.0	76.3	64.8	113.0	23.5	101.3
Test 5	24.8	24.3	23.0	21.3	34.3	7.5	30.5
Test 6	55.5	54.3	51.8	47.0	76.5	19.5	69.3
Test 7	29.0	29.3	27.5	28.0	34.3	21.0	32.8
Test 8	9.8	9.5	9.3	9.0	13.8	4.0	12.8
Test 9	69.3	68.3	61.8	65.3	90.8	32.0	84.3
Test 10	104.8	102.3	97.0	92.0	138.3	44.5	126.0
Test 11	107.3	105.5	96.3	101.5	152.5	41.5	141.3
Test 12	7.5	7.0	7.0	6.5	9.0	5.0	8.3
Test 13	26.3	25.3	24.0	23.5	35.3	11.0	32.8
Test 14	2.0	2.8	2.0	2.8	3.3	1.0	2.8
Test 15	3.0	2.8	2.8	2.3	3.5	1.5	3.5
Test 16	29.5	28.8	27.3	27.8	35.8	19.0	33.5

Table 1. Foxbase+/LAN Tests (Continued)

Flood Traffic (times in seconds)

	No Bridge	Hewlett-Packard HP 28674A	Magnalink 3000 HCB	RAD REB-20	Retix 4880	3Com NETBuilder IB/3000	Vitalink TransLAN 335
Test 1	61.0	69.8	68.3	170.3	72.8	150.0	80.3
Test 2	11.5	27.5	23.8	205.5	29.3	30.0	38.5
Test 3	16.5	50.3	38.3	433.8	53.8	53.5	72.3
Test 4	24.0	96.0	62.8	841.3	89.3	97.5	122.8
Test 5	7.5	28.8	21.5	248.0	27.0	406.0	36.5
Test 6	19.0	61.0	45.5	555.3	61.3	73.5	83.3
Test 7	20.5	28.0	27.8	113.3	30.5	30.5	37.0
Test 8	4.0	10.5	8.3	85.3	11.0	10.0	14.8
Test 9	32.0	65.8	64.5	520.8	74.5	75.0	100.5
Test 10	45.5	109.3	90.0	897.8	113.3	156.5	149.0
Test 11	41.5	105.8	99.0	1002.5	116.3	117.0	172.3
Test 12	5.5	7.5	7.0	64.3	8.0	8.0	9.5
Test 13	11.0	26.0	22.0	167.5	28.8	28.0	37.8
Test 14	1.0	2.3	2.5	15.3	2.3	2.0	3.3
Test 15	1.0	2.8	2.0	18.5	3.0	3.0	3.8
Test 16	19.0	28.5	28.0	136.5	30.8	30.5	37.8

11. Process 1,000-Transaction Batch
12. Select 1,000 Records on Indexed Field
13. Select 1,000 Records on Unindexed Field
14. Group and Subtotal 200-Record File and Report
15. Two-File Join and Report
16. Four-File Join and Report

Figure 10.
Throughput

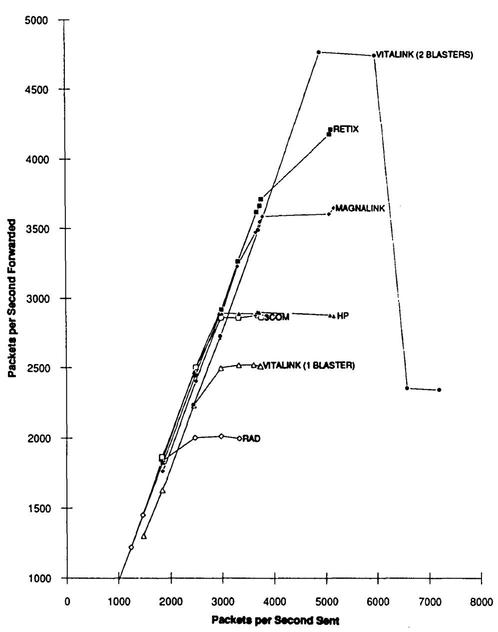

The 16-test suite is run without background traffic, then with four types of background traffic running concurrently on eight secondary workstations, four on each network segment.

Analysis

Except for the RAD, the bridges perform within a fairly tight range of times. The Magnalink is fastest overall for all traffic versions, followed closely by Hewlett-Packard. The Retix and 3Com times are very similar, and the Vitalink finishes just slightly behind because of its slower performance with Flood traffic.

The RAD is generally slower than the others for all traffic scenarios and especially with Flood traffic. Packet-per-second throughput limitations may have caused dropped packets, retransmission, and subsequently slower completion times with Flood traffic.

Throughput

Packets are transmitted from one secondary workstation to another across the bridges. All packets are the minimum Ethernet frame size (60 bytes plus a 4-byte frame check sequence and 8-byte preamble). Packets are sent at a range of speeds to determine the transmission rate at which the bridges can no longer forward packets fast enough to keep up.

Analysis

The transmitted frames are uniformly small and require less processing overhead than large frames; therefore, the test is biased in favor of bridges with short internal latencies for small packets and those with efficient buffering schemes that forward many frames in a short time. Bridges with small amounts of memory allocated for frame buffering may experience buffer overflow and consequently drop packets sooner than those with deep buffers. Dropped

Figure 11.
Internal Latency

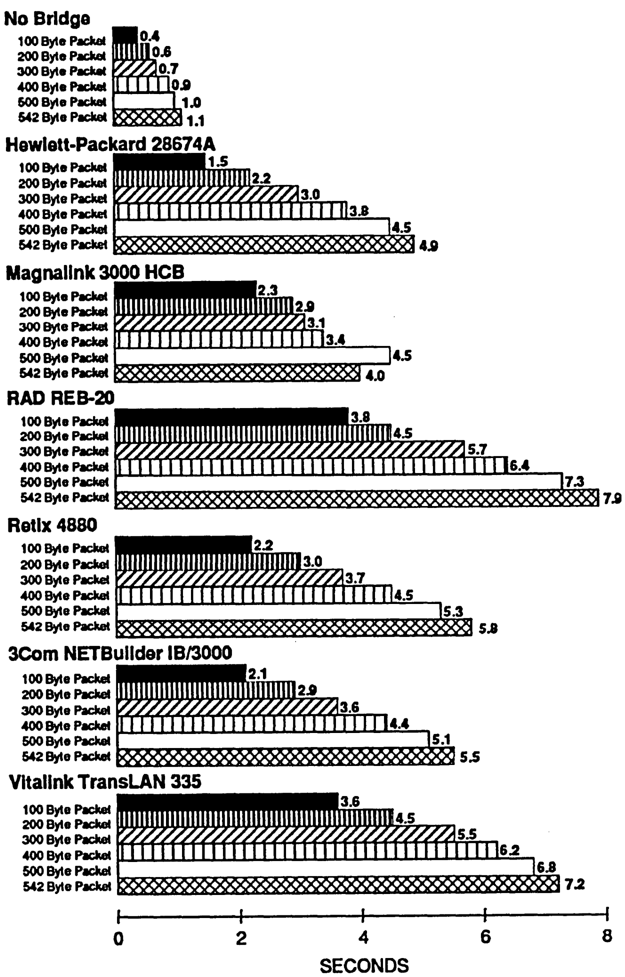

packets limit throughput, or the sustained rate at which a steady stream of packets can be forwarded. Bridges configured with two wide area links per bridge (Retix, 3Com, Vitalink, and RAD) have a clear advantage.

The Hewlett-Packard and 3Com both peak at about 2,900 packets per second, even though the 3Com has two wide area links and the Hewlett-Packard has one. The wide area link channel can handle 192,000 bytes per second, with a maximum of about 2,660 frames forwarded. The Hewlett-Packard achieves a level higher than this by stripping the 8-byte preamble and 4-byte frame check sequence from each frame prior to transmission. Each frame is encapsulated in an HDLC frame with a 2-byte header and 2-byte check sequence. Theoretically, the link bandwidth can forward three thousand 64-byte HDLC encapsulated frames. This type of procedure is not unique to the Hewlett-Packard bridge, but it explains how it achieves its throughput using a single T1 link.

The Vitalink achieves the highest throughput overall, with a marked drop at higher packet transmission rates. The Vitalink bridge pairs each support two wide area links and can effectively use both links (combined available bandwidth of 3.072M bps) only if traffic originates from more than one unique source address. The observed drop in throughput occurs intermittently at higher transmission rates and is typified by one wide area channel shutting down while the other remains active.

Internal Latency

The test measures the time required to transmit 3,000 frames between primary workstations across the bridge. Gross latency (in milliseconds per frame) is calculated only from runs in which no packets are dropped. The IPX send program will not send another packet until it receives an acknowledgement frame from the receive program. A time-out interval exceeding three seconds between transmissions causes a dropped packet to be logged.

Analysis

The RAD and Vitalink generally exhibit higher latencies or internal forwarding delays than the other bridges. Their performance results maintain a fairly linear relationship between packet size and latency for the range of packet sizes tested.

The Retix, 3Com, Hewlett-Packard, and Magnalink have overall shorter internal latencies, and all but the Magnalink maintain linear relationships between latency and packet size. The nonlinear character of the Magnalink curve may indicate the bridge's use of data compression to achieve high throughput levels.

Vendors

Hewlett-Packard Co.,
Business Computing Systems
19091 Pruneridge Avenue
Cupertino, CA 95014 (800) 752-0900

Magnalink Communications Corp.
P.O. Box 769, 63 Nahatan Street
Norwood, MA 02062 (617) 255-9400

RAD Network Devices
7711 Center Avenue, Suite 600
Huntington Beach, CA 92647 (714) 891-1964

Retix
2644 30th Street
Santa Monica, CA 90405-3009 (213) 399-2200

3Com Corp.
5400 Bayfront Plaza
Santa Clara, CA 95052 (408) 764-5000, (800) NET-3COM

Vitalink Communications Corp.
6607 Kaiser Drive
Fremont, CA 94555 (415) 794-1100

Characteristics

Table 2. Remote Ethernet Bridge Characteristics

	Filtering/Forwarding Rates	Address Table Size	Management Protocols	Memory (bytes)	Processor	Warranty and Support
Hewlett-Packard HP 28674A	Maximum 17,080/ 14,880 pps (64-byte packets)	512 entries	SNMP	Address Table: 512K; Frame Buffer: 256K	AMD 2900	1-year warranty covers parts, labor, and on-site service; telephone support ($45 per incident); service contracts available; extended warranty ($20/ month for 4-hour turnaround; $11/ month, 1-day turnaround)
Magnalink 3000 Series HCB	Maximum 14,000/ 3,800 pps (64-byte packets)	8,192 entries	SNMP (third-quarter 1991)	Address Table: 64K; Frame Buffer: 512K	68030	1-year warranty covers parts, labor, and return shipment; extended warranty ($1,100/year); telephone support
RAD REB-20	Maximum 10,000/ 2,600 pps (64-byte packets)	25,000 entries	None	Address Table: 20K; Frame Buffer: 240K	80186	1-year warranty covers parts, labor, and two-way shipment; extended warranty available at 10% of product price; toll-free telephone support
Retix High Performance 4880	Maximum 14,880/ 8,000 pps (64-byte packets)	16,000 entries	SNMP, MAP, TOP	Address Table: 96K; Frame Buffer: 500K	68020	1-year warranty covers parts, labor, and return shipment; extended warranty available; toll-free telephone support
3Com NET-Builder IB/3000	Maximum 19,000/ 10,000 pps (64-byte packets)	8,196 entries	SNMP	Address Table: 2M; Frame Buffer: 1M	68020	1-year warranty covers parts, labor, and return shipment; extended warranty available; hardware exchange for any hardware problem; toll-free telephone support ($1,995 covers 12 incidents)

Table 2. Remote Ethernet Bridge Characteristics (Continued)

	Filtering/Forwarding Rates	Address Table Size	Management Protocols	Memory (bytes)	Processor	Warranty and Support
Vitalink TransLAN 335	Maximum 14,880/ 7,000 pps (64-byte packets)	8,192 entries	None	Address Table: 164K; Frame Buffer: 4K	68020	1-year warranty covers parts, labor, and return shipment; extended warranty available; toll-free telephone support ($125/hour, 15-minute minimum) for customers not covered by maintenance

Table 3. Remote Ethernet Bridge Features

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NETBuilder IB/3000	Vitalink TransLAN 335
Physical Characteristics	0						
EIA rack-mount option	0	▲	▲	RA1	▲	—	▲
Weight (lb.)	0	6.01	12.5	RA1	17	12	16.5
Height (in.)	0	1.71	4.75	RA1	3.25	3.8	3.5
Depth (in.)	0	9.25	9.37	RA1	12.5	12.6	22.25
Length (in.)	0	16.75	17	RA1	17	16.2	17.2
Electrical/Environmental	0						
Power requirements (Watts)	0	60	61	200	50	155	100-130
Operating humidity (%), noncondensing	0	5-95	90	95	5-95	10-90	10-90
Operating temperature (degrees C)	0	0-50	0-50	40	0-50	5-40	5-40
FCC class (A or B)	0	A	A	A	A	A	A
Switchable 110/220 (115/230) V AC	0	▲	▲	▲	▲	▲	▲
Cooling fan	0	▲	▲	▲	▲	▲	▲
Power switch	0	—	—	▲	▲	▲	▲
Reset button	0	▲	▲	▲	—	▲	▲
DTE/DCE external switch	0	—	—	—	—	—	▲
Processor	0						
Processor type	0	AMD 2900	68030	80186	68020	68020	68020
Processor speed (MHz)	0	25	25	1	20	20	25
Embedded controller	0	▲	▲	—	▲	▲	▲
Memory	0						
Address table size (KB)	0	512	64	20	96	2MB	164
Frame buffer size (KB)	0	256	512	240	500	1MB	4

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Disk Subsystem	0						
Diskette drive	0	—	—	—	—	—	▲
Number of 3.5-inch diskette drives	0	0	0	0	1	1	2
Number of 5.25-inch diskette drives	0	0	0	1	0	0	0
Hard disk drive	0	—	—	▲	—	—	—
Hard disk capacity	0	—	—	40M	—	—	—
Console Connections	1						
RS-232-C	1	1	2	2	1	2	2
RJ-45	1	—	—	—	—	—	—
Other	1	—	—	—	—	—	—
LAN Connections	1						
BNC	1	1	1	0	1	2	0
AUI	1	1	1	1	1	2	1
Fiber optic	1	—	—	OPT	—	—	—
WAN Connections (Maximum)	1						
RS-422/449/530	1	0	1	4	2	2	8
RS-232-C	1	0	1	0	2	2	8
V.35	1	1	1	4	2	2	8
T1	1	1	0	4,OPT	2	0	8
Other	1	0	0	0	2,RE1	0	0
LAN Media Support	1						
Thick coaxial	1	▲	▲	▲	▲	▲	▲
Thin coaxial	1	▲	▲	▲	▲	▲	▲
Twisted pair	1	▲	▲	▲	—	—	▲
Fiber optic	1	—	—	▲	—	—	—
Diagnostics	1						
Power-up diagnostics	5	▲	▲	▲	▲	▲	▲
LED/LCD/Tone indicators	3	▲	▲	▲	▲	▲	▲
Console command line diagnostics	3	▲	▲	▲	▲	▲	▲
Downloadable diagnostics	1	▲	▲	▲	▲	▲	—
External LED/LCD Indicators	1						
Power (AC)	1	▲	▲	▲	▲	▲	▲
Alarm (12-Volt power failure)	1	—	▲	—	—	—	—
Self-test/diagnostic	1	▲	▲	▲	▲	▲	▲
Console	1	—	—	▲	▲	▲	▲
Network A collision	1	▲	MA1	RA2	▲	T1	▲
Network A transmit	1	▲	MA1	▲	▲	T1	▲
Network A receive	1	▲	MA1	▲	▲	T1	▲
Network B collision	1	▲	MA2	RA2	▲	T1	—
Network B transmit	1	▲	MA2	▲	▲	T1	—
Network B receive	1	▲	MA2	▲	▲	T1	—

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Management Software Connection	1						
Console (TTY; attached via serial connection)	1	▲	▲	▲	▲	▲	▲
In-band (attached via LAN)	1	HP1	▲	▲	▲	▲	▲
Modem support	1	▲	▲	▲	▲	▲	▲
Network Management Support	2						
Bridge management software	5	▲	▲	▲	▲	▲	▲
Bridge console commands	5	▲	▲	▲	▲	▲	▲
SNMP (Simple Network Management Protocol)	5	▲	MA3	—	▲	▲	—
CMIP (Common Management Information Protocol)	1	—	—	—	—	—	—
CMOT (Common Management Protocol over TCP/IP)	1	—	—	—	—	—	—
MAP (Manufacturing Automation Protocol)	1	—	—	—	▲	—	—
TOP (Technical Office Protocol)	1	—	—	—	▲	—	—
IEEE 802.3 Standard Support	1						
10BASE5 (standard/thick Ethernet)	2	▲	▲	▲	▲	▲	▲
10BASE2 (thin Ethernet)	2	▲	▲	▲	▲	▲	—
10BASE-T (unshielded twisted-pair Ethernet)	2	▲	▲	▲	—	—	▲
Fiber (Fiber Optic Internetwork Repeater Link)	2	▲	—	▲	—	—	—
Remote Communications	2						
Remote Links							
Leased lines	2	▲	▲	▲	▲	▲	▲
AT&T DDS	2	▲	▲	▲	—	▲	▲
Switched 56 (56KB dial-up)	2	▲	▲	▲	▲	▲	▲
T1	2	▲	▲	▲	▲	▲	▲
Fractional T1	2	▲	▲	▲	▲	▲	▲
Dial-up T1	1	▲	▲	▲	—	▲	▲
Other	1	—	MA4	—	—	T2	V1
Link Speeds							
Minimum link speed (K bps)	1	56	4.8	4.8	1.2	9.6	9.6
Maximum link speed (K bps)	1	2048	2048	2048	2048	7M	2048
Fallback dial-up	1	—	▲	▲	—	T3	▲
Fallback dial-up link speed (K bps)	1	—	56	4.8	—	7M	384

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Synchronous Communications Timing							
External clock	2	▲	▲	▲	▲	▲	▲
Internal clock	2	—	▲	▲	▲	▲	▲
Network (T1) timing	1	—	▲	▲	▲	—	▲
T1 extended superframe	1	—	MA5	▲	▲	—	▲
Load balancing among WAN ports (same bridge)	1	—	—	▲	▲	▲	▲
Load balancing among bridges	1	—	▲	▲	▲	—	V2
WAN port as backup	1	—	▲	—	▲	▲	▲
Other Remote Communications Features							
2-to-1 data compression	1	—	—	—	▲	—	—
4-to-1 data compression	1	—	▲	—	—	—	—
HDLC frame encapsulation	1	▲	▲	▲	▲	▲	▲
Sequenced frame transmission	1	▲	▲	▲	▲	▲	▲
High-Level Protocol Transparency	1						
TCP/IP	2	▲	▲	▲	▲	▲	▲
DECnet	2	▲	▲	▲	▲	▲	▲
LAT	2	▲	▲	▲	▲	▲	▲
XNS	2	▲	▲	▲	▲	▲	▲
ISO	2	▲	▲	▲	▲	▲	▲
Operating System Transparency	1						
Novell NetWare	2	▲	▲	▲	▲	▲	▲
3Com	2	▲	▲	▲	▲	▲	▲
Banyan	2	▲	▲	▲	▲	▲	▲
IEEE 802.3 compatibles	2	▲	▲	▲	▲	▲	▲
Performance (64-byte packets per second)*	0						
Filtering rate (maximum aggregate)	0	HP2	14K	10K	14,880	19,500	14,880
Forwarding rate, spanning-tree/learning enabled	0	12,953	3,800	RA3	8K	NA	7K
Forwarding rate, spanning-tree/learning disabled	0	14,880	3,800	2,600	8K	10K	7K
Throughput (M bps)	0	10	10	3	5	NA	8.192,V3
Source Address Database Configuration	3						
Automatic learning of source addresses	5	▲	▲	▲	▲	▲	▲
Maximum address table entries	1	512	8,192	25K	16K	8,196	8,192
Maximum remote address entries	1	512	1,024	25K	16K	8,196	8,192
Maximum addresses in nonvolatile RAM	2	512	1K	100	256	512	SYS

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Source Address Database Configuration (Continued)							
Automatic aging of source addresses	2	▲	▲	▲	▲	▲	▲
Accelerated aging upon spanning-tree change	1	▲	▲	—	▲	—	▲
Manually add/delete address records	2	▲	▲	▲	▲	▲	▲
Exempt address record from aging	1	▲	▲	▲	▲	▲	▲
Lock out changes to addresses	1	▲	▲	▲	▲	▲	▲
Auto correct subnet assignment	1	—	—	—	—	▲	▲
Save addresses in nonvolatile RAM	2	—	▲	▲	▲	▲	—
Symbolic names for address records	1	▲	BR	▲	▲	—	—
Add/delete range table address	1	▲	▲	—	—	—	—
Save/retrieve range table configurations	1	▲	▲	—	—	—	—
Set aging time interval	1	▲	▲	▲	▲	▲	▲
Set symbolic names on/off	1	—	—	▲	—	—	—
Set address dispositions (flood, forward, discard)	2	▲	▲	▲	▲	▲	▲
Set upper/lower bounds for range table entries	1	—	▲	—	—	—	▲
Set dispositions for range table entries	1	—	▲	—	—	—	▲
Source Address Database Statistics	2						
All source address records	2	▲	▲	▲	▲	▲	▲
Restricted source address records	2	▲	▲	▲	▲	▲	▲
Static source address records	2	▲	▲	▲	▲	▲	▲
Source address records in nonvolatile RAM	2	▲	▲	▲	▲	▲	—
Protocol-/Operating System-Specific Filtering	2						
AppleTalk	1	▲	▲	▲	—	▲	▲
ARP	1	▲	▲	▲	—	▲	▲
Bridge ID	1	▲	▲	▲	—	▲	▲
IP	1	▲	▲	▲	—	▲	▲
DECnet	1	▲	▲	▲	—	▲	▲
LAT	1	▲	▲	▲	—	▲	▲
NetWare	1	▲	▲	▲	—	▲	▲
REVARP	1	▲	▲	▲	—	▲	▲
TCP/IP	1	▲	▲	▲	—	▲	▲

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Protocol-/Operating System-Specific Filtering (Continued)							
TOP	1	▲	▲	▲	—	▲	▲
V2	1	▲	▲	▲	—	▲	▲
XNS	1	▲	▲	▲	—	▲	▲
User-defined protocol filtering	1	—	—	—	—	▲	▲
Routing Methods	1						
Source address	1	▲	▲	▲	—	▲	▲
802.1 spanning-tree protocol	1	▲	▲	▲	▲	▲	▲
Automatic cable loop avoidance	1	▲	▲	▲	▲	▲	—
Filtering	2						
Filter frame types	1	▲	▲	▲	▲	▲	▲
Pass (always forward) by packet type	1	▲	▲	▲	▲	▲	▲
Filter frames by source address	1	▲	▲	▲	▲	▲	▲
Filter frames by destination address	1	▲	▲	▲	▲	▲	▲
Filter frames by packet size	1	—	▲	▲	—	▲	▲
Enable/disable filters	1	▲	▲	▲	▲	▲	▲
Build/maintain filter tables	1	▲	▲	▲	▲	▲	▲
Filter frames by address range	1	—	▲	▲	—	▲	▲
Forwarding	2						
Auto forward unknown destination address	2	▲	▲	▲	▲	▲	▲
Auto forward broadcast frames	2	▲	▲	▲	▲	▲	▲
Group address forwarding	1	—	▲	▲	—	▲	▲
Individual address forwarding	2	▲	▲	▲	▲	▲	▲
Enable/disable forwarding of multicast frames	1	▲	▲	▲	▲	▲	▲
Forward frames within address range	1	▲	▲	▲	—	▲	▲
Administration and Configuration	3						
Modify serial comm settings for console connection	1	▲	▲	▲	▲	▲	▲
Console command line Help	1	▲	▲	▲	▲	▲	▲
Bridge on-/off-line	1	▲	▲	▲	—	▲	▲
List devices on extended net by manufacturer	1	▲	▲	▲	—	—	▲
Update reset log	1	▲	▲	▲	—	▲	▲

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Administration and Configuration (Continued)							
Software reset with updated configuration	1	▲	▲	▲	▲	▲	▲
Hardware reset with updated configuration	1	▲	▲	▲	▲	▲	▲
Software reset with default configuration	1	—	▲	▲	▲	▲	▲
Hardware reset with default configuration	1	—	▲	▲	▲	▲	▲
Report Network Management Center (NMC) IP address	1	▲	—	—	—	▲	—
Console control over multiple bridges on extended net	1	▲	▲	▲	▲	▲	▲
Save parameter/modifier settings in nonvolatile RAM	1	▲	▲	▲	▲	DISK	—
Modular software cartridge	1	—	—	—	▲	▲	—
Reset all bridge statistics	1	—	▲	▲	▲	▲	▲
Initialize bridge start-up/diagnostic cycle	1	▲	▲	▲	▲	▲	▲
Logout of management session	1	▲	▲	▲	▲	▲	▲
Integrated text editor	1	▲	—	▲	—	—	▲
Save/retrieve profile configuration	1	▲	—	▲	▲	▲	▲
Monitor multiple bridges	1	▲	▲	▲	▲	▲	▲
Menu-driven management	1	▲	▲	▲	▲	▲	▲
Management software included	1	—	▲	—	▲	▲	▲
Help function	1	▲	▲	▲	▲	▲	▲
Set console command line prompt	1	—	▲	▲	▲	▲	▲
Set bridge name	1	▲	▲	▲	▲	▲	▲
Set network segment names	1	—	—	▲	▲	▲	▲
Set bridge IP address	1	▲	—	—	▲	▲	▲
Set IP NET mask for IP routing	1	▲	—	—	▲	▲	▲
Set IP address of IP router	1	▲	—	—	—	▲	▲
Set standard/extended packet size	1	—	—	▲	—	—	▲
Name downloadable bridge software file	1	▲	—	▲	▲	▲	▲
Set console event reporting level	1	—	▲	—	▲	—	—
Forward operation mode	1	▲	▲	▲	▲	▲	▲
Filter operation mode	1	▲	▲	▲	▲	▲	▲
Pass All operation mode	1	—	▲	▲	▲	▲	▲
Set login message	1	—	—	—	—	▲	—

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Administration and Configuration (Continued)							
Statistics event threshold options	1	▲	▲	▲	▲	—	—
Statistics gathering options	1	▲	▲	▲	▲	—	—
No-activity time limit to reset transmitter chip	1	—	▲	▲	▲	—	▲
No-activity time limit to reset receiver chip	1	—	▲	▲	▲	—	▲
Set forwarding delay	1	—	—	—	▲	—	▲
Source address learning mode on/off	1	—	▲	—	▲	▲	▲
Source address record aging on/off	1	▲	▲	—	▲	▲	▲
Set port state (Listen, Learn, Forward, Block, Disable)	1	▲	▲	—	▲	▲	▲
Set port priority level	1	▲	▲	—	▲	▲	▲
Display Statistics	2						
Bridge configuration/status	1	▲	▲	▲	▲	▲	▲
Specific bridge parameters	1	▲	▲	▲	▲	▲	▲
Bridge name	1	▲	▲	▲	▲	▲	▲
Reset log	1	▲	▲	▲	▲	—	▲
Downloadable bridge software filename	1	—	—	▲	▲	—	▲
Console event reporting level	1	—	▲	▲	—	—	—
Bridge operation mode	1	▲	▲	▲	▲	▲	▲
Names of attached network segments	1	—	—	▲	▲	▲	▲
Login message	1	—	▲	—	—	▲	—
Ethernet bridge address	1	▲	▲	▲	▲	▲	▲
All bridges on network	1	▲	▲	▲	▲	—	▲
Packets transmitted/received	1	▲	▲	▲	▲	▲	▲
CRC and alignment errors	1	▲	▲	▲	▲	▲	▲
No-resource errors	1	▲	▲	—	▲	▲	▲
Oversize/undersize packets	1	▲	▲	▲	▲	▲	▲
Packet transmission failures	1	▲	▲	▲	▲	▲	▲
Packet transmission collisions	1	▲	▲	—	▲	▲	▲
Packets queued for transmission	1	▲	—	—	—	▲	▲
SNMP packets broadcast	1	—	—	—	—	▲	▲
Valid SNMP packets transmitted	1	—	—	—	▲	▲	▲

Table 3. Remote Ethernet Bridge Features (Continued)

	Weight	Hewlett-Packard HP 28674A	Magnalink 3000 Series HCBG	RAD REB-20	Retix 4880	3Com NET-Builder IB/3000	Vitalink TransLAN 335
Flags	1						
Error severity flags	1	—	—	—	—	—	—
Forwarding flag for unknown group addresses	1	—	—	▲	—	—	—
Forwarding flag for unknown individual addresses	1	—	—	▲	▲	—	—
Origin filter learning flag for source addresses	1	—	▲	—	▲	—	—
Restricted learning flag for destination addresses	1	—	▲	—	—	▲	—
Static learning flag for addresses	1	—	▲	▲	▲	—	—
Spanning-Tree Protocol Settings	1						
Standby operation mode (backup)	1	▲	▲	—	▲	▲	▲
Root priority level	1	▲	▲	—	▲	▲	▲
Hello time	1	▲	▲	—	▲	▲	▲
Bad Hello limit	1	—	▲	—	—	▲	▲
New tree time	1	▲	▲	—	▲	▲	▲
Hello maximum age	1	▲	▲	—	▲	▲	▲
Secure mode	1	▲	—	—	—	—	▲
Enable spanning-tree protocol	1	▲	▲	—	▲	▲	▲
Disable spanning-tree protocol	1	▲	▲	—	▲	▲	▲
Set path cost per subnet	1	▲	▲	—	▲	▲	▲

▲—Yes, has feature.
BR—Bridges only.
DISK—Saved to disk; through console using a workstation.
HP1—Requires Hewlett-Packard OpenView Bridge Manager.
HP2—17,080 LAN and WAN maximums.
MA1—LAN port.
MA2—WAN port.
MA3—Expected third-quarter 1991.
MA4—V.32.
MA5—Per CSU/DSU.
NA—Not available.
OPT—Optional.
RA1—Depends on PC platform used.
RA2—Low, Medium, and High traffic indicators.
RA3—Spanning tree not available.
RE1—X.21.
SYS—System addresses only.
T1—One LED activity indicator.
T2—Fractional T3 to 7M bps.
T3—Via NCS/AT.
V1—Private networks.
V2—Through filtering.
V3—1,518-byte frames.
**Per manufacturers' specifications.* ■