Hewlett-Packard Series 200, 400, and 600 Routers

HP Routing Services and Applications

Publication Number
5962-8770E

Edition 1, July 1994

Printed in Singapore

Product Numbers and Software Version

This guide provides information for Hewlett-Packard routers running software with the following version numbers:

A.08 series

B.08 series

C.08 series

Earlier and later software versions may operate differently than described in this manual.

Warranty

The information contained in this guide is entirely unwarranted.

Hewlett-Packard
8000 Foothills Boulevard
Roseville, California 95747-6588
916-786-8000

Contents

1 Product Notes

Features of HP Routers 1-3

Architecture and Technology 1-21

Branch Office Routing 2-35

2 Routing Services Notes

Bridging Service 2-3

Internet Protocol Routing Service 2-55

Novell IPX Routing Service 2-95

AppleTalk Phase 2 Routing Service 2-107

DECnet Routing Service 2-115

A Primer on HP Probe 2-125

Data Compression for WAN Links 3-133

3 Application Notes and Case Studies

Improving Network Availability 3-3

ISDN Wide Area Network Design:
Dry Creek Joint Elem. School District 3-31

Shining a Light on FDDI 3-41

Using Synchronous Pass-Through to Consolidate
Synchronous Traffic 3-57

Routing with OSPF 3-69

Linking Up with Frame Relay 3-85

Frame Relay Network Design: Fleet Call, Inc. 3-105

1

Product Notes

Product Notes

- Features of HP Routers
- Architecture and Technology
- Branch Office Routing

Features of HP Routers

Each HP router is a multiprotocol router and multiport bridge. The router is used to form an internetwork using local area connections (such as Ethernet, IEEE 802.3, token ring, and FDDI) and wide area connections (such as leased lines, X.25, frame relay, SMDS, and ISDN services).

This product note summarizes the features of HP's AdvanceStack line of routers in general. Each specific router product may not support all of the features described here. For cost-effective branch routing, for example, some routers in the 400 series have a fixed selection of ports, and routers in the 200 series have a partial selection of routing services. And for a central site router, the HP Router 650, port modules allow you to select the types and number of port connections needed. For the features available and unavailable on specific routers, see the *HP Network Connectivity Product Catalog*, or the installation manual and release notes for your specific router product.

Routing Services and Protocols

The standard configuration of HP routers supports concurrent operation of five popular routing services: Internet Protocol (the TCP/IP protocol suite), Novell IPX, AppleTalk Phase 2, DECnet Phase IV, and Xerox XNS. For packets that cannot be routed, the router can function as a learning bridge, source-routing bridge, or translational bridge, and can use the IEEE 802.1 Spanning Tree Protocol.

Each routing service may be enabled independently of the others, and it builds and maintains its routing database independently of the others. Likewise, the bridging service may be enabled independently of any of the routing services; if not enabled, then the router will drop packets with protocol types not enabled or supported (such as IBM SNA or DEC LAT). Because of the independent routing services, the HP router is often referred to as an IP

router, a DECnet router, a bridge, and so on, when the associated service is enabled.

On HP routers, routing services should be used when possible for their bandwidth-conserving value, and they must be used for links between dissimilar physical-layer and data-link-layer operations. Bridging is usually required as well for unsupported protocols and for such services as synchronous pass-through.

Bridging

With bridging enabled, an HP router provides protocol-transparent bridging between similar LAN types, such as Ethernet to Ethernet. It learns the location of nodes on the network—based on their data-link-layer addresses—and builds a dynamic bridging table for use in deciding whether or not to forward an incoming frame. Frames may be bridged to other LANs that are local or located at a remote site over wide area connections. It is possible to use the HP router solely as a bridge by not enabling any of the routing services.

As an Ethernet/IEEE 802.3 transparent bridge (or learning bridge), an HP router supports a variety of filtering options at the MAC layer, such as source and destination address and protocol type. The IEEE Spanning Tree Protocol (STP) is available for management of bridged networks with mesh topologies. STP allows a bridged network to have redundant paths. In the event of a primary link failure, a backup link takes over, thereby ensuring continued data transmission between all reachable network segments.

Translational bridging is used between FDDI and either Ethernet or IEEE 802.3 networks. Source-route translational bridging is used for bridging between token ring (IEEE 802.5) and Ethernet or IEEE 802.3 networks. These addressing schemes differ at the data-link layer, specifically at the MAC layer. For more information about translational bridging, refer to the “Bridging Service” note and to the “Shining a Light on FDDI” application note in this manual.

Source-routing bridging is used to connect token rings that contain systems communicating with non-routable protocols such as IBM SNA or NetBIOS. Thus, when an HP router is configured for source-routing bridging, it can be used in any application that would otherwise be performed by an IBM Source Route Bridge. Transparent (also known as learning) bridging is always enabled with the bridging software, so that when source-routing bridging is performed concurrently, the HP router is known as a source-routing/transparent (SRT) bridge.

Notes

However, HP routers do not *transparently* bridge from a token ring to a token ring (only between Ethernet/802.3 LANs); source-routing bridging is used for this function. Also note that the source-routing bridging function does not provide communication between a source-routing system on one ring and a non-source-routing system on another ring.

For details on the various types of bridging, refer to the "Bridging Service" note later in this manual.

HP routers provide source-routing end-node support for the IP, IPX, AppleTalk, and XNS routing services (described in the following sections). These routing services are users of the underlying physical and data-link layers, which allow them to route packets in token ring and mixed-media environments.

IP Routing

The Internet Protocol (IP) is the network layer of the TCP/IP de facto standard. It is supported by most major computer vendors and used extensively in large networks such as the "Internet".

Generally, each corporation or similar entity sets up an internal network composed of a number of routers and computers. Within this network, IP allows the creation of subnets. Subnets usually reflect the organizational and geographic structure of the entity. Communication between subnets can optionally be enabled or disabled based on subnet address. This capability simplifies traffic control and enhances network security.

Within the entity's internal network, IP routers exchange network information using an interior gateway protocol (IGP). HP routers support two IGPs: Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) protocol.

RIP is widely accepted as a standard router-to-router protocol for IP networks. RIP selects network routes based on the lowest number of hops required to traverse the network between subnets. In addition to RIP's standard features, HP routers allow users to adjust the hop count for a particular network link to reflect lower link speed, higher delay, or other factors. This feature allows RIP to be tuned to avoid network bottlenecks.

OSPF is another IGP standard that is better suited to larger, more complex networks than RIP. OSPF provides more security with an option for password authentication on routing update messages. OSPF increases network efficiency by using multicast addresses for update messages, sending

incremental routing updates, and supporting a hierarchical update structure (using areas).

RIP and OSPF are described in more detail in the "Internet Protocol Routing Service" note later in this manual.

For connections to external IP networks, HP routers support the standard Exterior Gateway Protocol (EGP). EGP provides secure connections through designated routers to external networks.

The IP router software for HP routers provides full support for transport-layer protocols such as Transmission Control Protocol (TCP), the Internet Control Message Protocol (ICMP), and the User Datagram Protocol (UDP). Both ARPA/Berkeley and HP NS upper-layer services are supported, including the ARP (Address Resolution Protocol) and HP Probe address resolution protocols. For more information about HP Probe, refer to the note "A Primer on HP Probe" in this manual.

Novell IPX

HP routers support the Internet Packet Exchange (IPX) and Sequenced Packet Exchange (SPX) protocols used by Novell's NetWare. The IPX routing service integrates NetWare-based PC networks into the enterprise LAN. The IPX routing service on HP routers is compatible with network nodes that use Novell NetWare version 2.15 or later. HP routers can interoperate with routers using Netware and with Novell routers—by negotiating data-link characteristics when establishing a WAN link (over PPP, X.25, or frame relay). IPXWAN Version 2 is supported. For more information about IPX, refer to the "Novell IPX Routing Service" note in this manual.

AppleTalk

AppleTalk is the protocol suite developed by Apple Computer, Inc., for LAN communications between Macintosh personal computers. HP routers support AppleTalk Phase 2, which enhances the original AppleTalk protocol suite (Phase 1) by expanding the addressing and naming capabilities previously available. The AppleTalk Phase 2 standard also provides support for Ethernet and token ring networks in addition to LocalTalk, Apple's proprietary data-link-layer protocol.

HP routers support AppleTalk Phase 2 routing over Ethernet links or synchronous WAN links. LocalTalk links to HP routers are not available. AppleTalk Phase 1 routing is not supported; however, it can be relayed using the bridging function of the router. For more information, refer to the "AppleTalk Phase 2 Routing Service" note in this manual.

DECnet Phase IV

The DECnet routing service for HP routers supports the DECnet Phase IV routing standard. As specified by this standard, large networks are subdivided into "areas". Typically, areas are composed of one or more geographically separate LANs. Each router is part of one of these areas; there can be multiple routers within an area. Routing within an area is called level 1 routing and routing between areas is called level 2 routing. HP routers support both level 1 and level 2 routing.

HP routers can interoperate over Ethernet with routers from DEC; however, direct interoperability over serial links (WAN links) is not supported. For more information, refer to the "DECnet Routing Service" note in this manual.

Note

The DECnet routing service does not support routing on token ring (IEEE 802.5) networks.

Xerox XNS

The XNS (Xerox Network Systems) routing service on HP routers is defined by the Xerox Graybook standard for the XNS protocol stack. Routing is based on the Internet Datagram Protocol (IDP). The Routing Information Protocol (RIP), Error Protocol, and Echo Protocol used in XNS are also supported.

Proprietary adaptations of XNS used by vendors such as 3Com and Ungermann-Bass are not currently supported. Apollo Domain routing, which is also based on Xerox XNS, is not currently supported.

Table 1. Routing Service Protocols on HP Routers

Routing Service ("Routable" Protocol)	Network-Layer Communication Protocol	Network-Layer Routing Protocol	Data-Link-Layer Address Translation Protocol (to get station address of node)
AppleTalk Phase 2 routing	AppleTalk Datagram Delivery Protocol (DDP)	Routing Table Maintenance Protocol (RTMP)	AppleTalk Address Resolution Protocol (AARP)
DECnet Phase IV routing	DECnet routing protocol (DRP)	DECnet routing protocol (DRP)	Not needed (station address assigned and used as network-layer address)
IP routing (TCP/IP protocol suite)	Internet Protocol (IP)	Exterior Gateway Protocol (EGP) Interior gateway protocols: ▪ Routing Information Protocol (RIP) ▪ Open Shortest Path First (OSPF) Static routes	HP Probe Address Resolution Protocol (ARP) DDN (Defense Data Network) address resolution algorithm PDN (public data network) RFC-877 compliant address resolution
Novell IPX routing	Novell Internetwork Packet Exchange (IPX)	Routing Information Protocol(RIP) Static routes	Not needed (station address used as network-layer address)
Xerox XNS routing	Internet Datagram Protocol (IDP)	Routing Information Protocol (RIP) Static routes	Not needed (station address used as network-layer address)
Bridging for non-routable protocols	n/a	n/a	Not needed (the bridge uses only the data-link-layer station address)

Security and Traffic Filtering

HP routers provide network security, isolation, and traffic control with powerful user-configurable traffic filtering based on protocol (packet type), service type (port, socket), network/subnet address, and network node addresses (users).

AdvanceStack Router Hardware

Series 600: Modular Routers

The HP Router 650 is used as a high-speed central-site router. It is a compact table or rack-mountable chassis with four slots for LAN and/or WAN interface card modules. Most of the interface module products each have four network ports, and can be installed in any combination. See the *HP Network Connectivity Product Catalog* for the specific interface modules and port combinations available. The types of ports available on these modules are described in following sections of this note.

HP's Ethernet/IEEE 802.3 interface module product itself has slots that accept either AUI, BNC, 10Base-FL (fiber-optic), or twisted-pair type of recessed miniature transceivers. (See "Ethernet Ports" on the next page for details.) A fifth slot in the chassis is occupied by the built-in routing engine PCA; it includes a console port and router resetting controls. The chassis also contains a module of variable-speed fans and one auto-ranging power supply. A second power supply can be added for load sharing and redundancy. All interfaces, the redundant power supply, and the fan module can be replaced online without interrupting network connections and components that are still operating ("hot swap"). An LED PCA provides multiple status indicators for each port on each card, for overall chassis status, and for the fan and redundant power supply. For more information, see the *Installation Guide* for the HP Router 650, and in this manual, see the "Architecture and Technology" note.

Series 200 and 400: Fixed-Configuration Routers

Routers such as the HP Router FR and HP Router SR, and several others, are smaller and self-contained, each with a fixed combination of two or more ports for different applications as regional and branch routers in your network topology. Some of these stackable and rackable routers use one 1¾-inch rack space; others use two 1¾-inch rack spaces. The types of ports available on these routers are described in the following sections of this note. See the *HP Network Connectivity Product Catalog* for the specific routers and port combinations available.

Warranty

- 3-year on-site warranty

LAN Media Connections

The following types of ports are available on HP routers. See the *HP Network Connectivity Product Catalog* for the specific routers and their port combinations.

Ethernet Ports

Ethernet/IEEE 802.3 LAN ports use the following types of connectors:

- Twisted-pair RJ-45 connector, for IEEE 802.3 Type 10Base-T unshielded twisted-pair cable
- BNC connector, for a IEEE 802.3 Type 10Base2 thin coaxial cable LAN (also known as ThinLAN) to be attached with a BNC "T" connector
- Fiber-optic transmit and receive ST-type connectors, for IEEE 802.3 Type 10Base-FL (or FOIRL) optical cable
- AUI connector, for an external Ethernet/IEEE 802.3-compatible transceiver used to connect to any of the other LAN media listed above this item, plus the following medium:
- Thick coaxial cable, IEEE 802.3 Type 10Base5

On an HP Router 650, the Ethernet/802.3 ports are supplied on replaceable recessed transceiver modules. All of the Ethernet/802.3 ports on the series 200 and 400 routers include an AUI connector for attachment of an external transceiver, and some also include a BNC connector that you can use instead of the AUI port.

Token Ring Ports

Each token ring port has a 9-pin female connector. This is used to connect an IEEE 802.5 token ring LAN using a cable with a 9-pin male D-connector on one end and a medium interface connector (MIC for token rings) on the other end of the cable. The MIC must conform to IEEE 802.5 specifications for connection to a trunk coupling unit (TCU).

FDDI Ports

An FDDI port is used to connect an FDDI ring to the router using a class A dual attachment station (DAS) or class B single attachment station (SAS) or dual homing. An optical bypass switch can be connected between the router and the DAS or SAS. The port has two FSD (fixed shroud duplex) connectors, MIC A and MIC B (media interface connectors for FDDI rings). Both A and B are used to attach a DAS; either A or B is used to attach a SAS. The separate optical bypass connector is used, in addition, to attach the optional-bypass switch (dual-switch module). Multimode fiber is required. The wavelength used is 1300 nanoseconds.

WAN Connections and Services

The wide-area synchronous ports on HP routers do the following:

- Provide access to frame relay, SMDS, and X.25 packet-switching networks.
- Provide access to switched circuits, including the public switched telephone network (PSTN or POTS), switched 56-Kbit/second, switched fractional T1, and ISDN networks.
- Support link-terminating equipment, including modems (on the PSTN), ISDN manual terminal adapters (modems in DTR mode), V.25 terminal adapters for ISDN, and CSU/DSUs for the switched digital network.
- Support data compression over all point-to-point links. Whether data compresssion is enabled is automatically negotiated between the routers.
- Allow load sharing between multiple WAN links of equal bandwidth for the best use of leased lines between remote routers. Other bandwidth management features are traffic prioritization and latency control.
- Support synchronous WAN connections using the HP universal interface ports with the appropriate 5-meter interface cable (V.35, X.21, RS-232 or V.24/V.28, RS-422/449 or V.36). HP routers automatically sense which specific interface cable is connected. Wide area link speeds up to 2.048 megabits per second are possible.

WAN Ports

On the synchronous ports, wide area network connections are established by connecting WAN-link-terminating equipment. This equipment, which provides timing signals to the router's WAN interface, may be typically a data service unit/channel service unit (DSU/CSU), a modem, or an ISDN terminal adapter (using the V.25 bis standard protocol). The WAN port connectors on the rear of the routers are 62-pin high-density female D-subminiature connectors. When ordering the router, one of four types of personality cables is selected for the WAN ports—to adapt the 62-pin interface to one of the four WAN types: RS-232 or V.24/V.28 (with a 25-pin male connector), RS-422/449 or V.36 (with a 37-pin male connector), V.35 (with a 34-pin male connector), and X.21 (with a 15-pin male connector).

X.25 WAN Links

AN HP router can transfer data through a public or private X.25 packet-switching network to another HP router. In addition, IP traffic from the router can be sent through an X.25 network directly to any computer system that supports the standard for IP encapsulation in X.25 (RFC 877). HP routers' X.25 services comply with the CCITT 1984 recommendation for X.25 and with the U.S. Defense Data Network (DDN) X.25 standard. Switched virtual circuits (SVCs), statically configured SVCs, and permanent virtual circuits (PVCs) are supported (up to a maximum of 255 virtual circuits per interface module in the Router 650). Variable packet sizes from 16 to 2048 bytes are handled, and line speeds up to 64 kilobits per second are handled.

Other WAN Services

- HP routers can use a proprietary point-to-point WAN protocol between themselves. Or, for interoperation in multivendor environments, the industry-standard Point-to-Point Protocol (PPP) can be used. Dynamic link configuration allows bridging to be enabled automatically over leased lines, using the HP-proprietary protocol, between two HP routers or between an HP Remote Bridge and an HP router.
- Synchronous traffic (SDLC, HDLC, LAP-B) can be integrated with routing traffic over the WAN links using the synchronous pass-through feature. See the application note, “Using Synchronous Pass-Through to Consolidate Synchronous Traffic”, later in this document.
- Dial on demand: Using ISDN switched circuits, the router can be configured to have a V.25 bis terminal adapter automatically dial different phone numbers for primary and backup circuits, or configured to dial a different number for each destination IP address. Some more information is found in the application note “Improving Network Availability” in this manual.
- Various features that help control WAN costs—by improving performance, improving robustness, and planning and controlling WAN capacity—are described in the application notes, “Improving Network Availability” and “Branch Office Routing”, later in this manual. Bridged traffic can be prioritized over routed traffic. Minimum bandwidth can be reserved for lower priority traffic. Transmit buffers can be tuned to control latency.

"Instant On" and SmartBoot

Implementing an internetwork with routers should not require that numerous network administrators be trained and staffed at every site in an organization. Ideally, routers should be installable and manageable by centralized MIS administrators remotely. The simplified solution from HP requires only that branch personnel place the router on a shelf, connect communication cables, and apply power. The router begins network service automatically, and any further configuration and administration can be done by central site personnel. HP's "Instant On" features include the following.

- All configuration is done through software, and can be done remotely; no hardware configuration settings are involved.
- Software is preloaded; no disks to insert; no need to download the operating system.
- WAN interface type is automatically detected from the type of cable attached; no switches or configuration needed.
- WAN link configuration of leased lines to another HP router or bridge is auto-detected and negotiated; no configuration for data-link layer is needed.
- SmartBoot: Branch office routers acquire their configuration automatically when placed in service; no expert personnel required at the branch site, only at the central site.
- Detailed configuration at a console attached to the router not required for bridging and/or routing to begin.

Branch Routing with SmartBoot

Once the cables are attached, HP routers automatically establish network connections with the LANs and with other corporate or regional HP routers (and HP Remote Bridges) on WAN links. Dynamic link configuration is used for WAN links; preconfiguration is not required unless there are special link requirements such as X.25, frame relay, or a phone number to automatically dial. All links will bridge traffic at first.

Then, the router automatically acquires its configuration over the network, allowing it to perform routing. Two basic scenarios for this autoconfiguration over the network are the following.

1. The administrator of the router at the central/regional site has used Quick Remote—a component of SmartBoot included on HP routers—to create and store a configuration file for each of the remote routers using its WAN links. The configuration file includes the most basic information—such as network addresses and which of the services (bridging, IP routing, and IPX routing) are to be enabled—for no more than one WAN port and one LAN port on the branch router. The branch router, when in the factory default state, uses the Bootp protocol to request its configuration and uses it to boot up to begin routing. No Bootp file server—just the other HP router—is needed. (Quick Remote is described in the next section of this note.)

   Branch router configurations can be further customized on a remote console over the network (Telnet) or out of band over the RS-232 console connection (using a modem, for instance).

Note

You don't need a specific router from HP for branch routing.
Any HP router except the HP Router PR can be the central router.
Any router except the HP Router 650 can be the remote router.

2. A network administrator has stored a configuration file for the HP router on any Bootp server on the network. HP provides applications that specify the most basic routing information and others that create fully customized configuration files, for each specific router in the network. Or, a configuration previously used on the network can be transferred (and individualized) and stored on a Bootp file server for use by another router. This scenario does not require another HP router.

   The branch router, from its factory default state, then uses the Bootp protocol to request its configuration. If it receives a reply from a Bootp server, with the filename of a configuration file reserved for that router's IP address, the router will automatically use the TFTP protocol (the Trivial File Transfer Protocol) to download the file and use it to boot up and begin routing. Branch router configurations can be further customized at any time.

Dynamic WAN Link Configuration

HP routers auto-detect and negotiate the following when connected to an operational link:

- Device type of the HP router or HP Remote Bridge attached
- LLC protocol type (quality of service)
- Clocking and HDLC device address DCE or DTE
- Compression setting
- Link speed to automatically configure transmit queues (for reduced latency)

This works with WAN link types (such as leased lines) that require no specific configuration (that is, excluding such features as X.25, frame relay, or a phone number to automatically dial). Also, multivendor connections usually require specific configuration before starting up.

Easy Distribution of Software and Configurations

The operating system software is preloaded and stored in flash EEPROM. There are no disks to insert. Software can be updated over the network.

TFTP (the Trivial File Transfer Protocol) can be used to push or pull either the operating system or the configuration file through the network. This can be done from one centralized facility on the network. Each router can be configured with a list of IP addresses from which it will accept these files.

On a PC connected as a console, the Zmodem protocol can also be used to store and distribute router configurations, as well as the output of router status, event logs, routing tables, statistics, and configuration displays. With operating system updates, HP provides an update utility for using the console.

Console (RS-232) Port

The router has an RS-232 port specifically for out-of-band console access, using an ANSI or VT100 terminal or a PC running an ANSI or VT100 terminal emulation program, and optionally a modem. Information on terminals and modems and cables can be found in the installation manual for your router. Console access is secured with both a manager-level and a user-level password. The same console access is available in band, using Telnet (remote terminal session).

Many router configuration and management functions can be performed on a terminal or PC connected through the router's RS-232 asynchronous serial console port. The following is a summary of the functions you can perform:

- Set date and time and passwords.
- Configure links and routing and bridging services.
- View the current configuration.
- Manually enable and disable configured links and routing and bridging services.
- Review logged events to monitor the links and services being established.
- View online network traffic statistics.
- View routing and bridging tables.
- Access MIB variables.
- Establish a remote console or terminal session on another node using Telnet.
- Transfer operating system, configuration, and console screen data to another node using the Trivial File Transfer Protocol (TFTP).
- Download new router operating system enhancements.
- Execute data-link-layer and network-layer diagnostic tests.

Configuration

Routers are shipped from HP with a default configuration that allows them to come up in their attached networks as bridges. To perform routing with the various available protocols (such as IP, IPX), each interface will need at least some protocol- or network-specific configuration information. The routers used in branch offices can obtain from another connected router enough network-specific information to come up as routers. This SmartBoot feature is described above. For further customization and tuning of the router's configuration, you can use one or both of the local configuration utilities available on the router's console, Quick Configuration and Configuration Editor.

Quick Configuration Router configuration in a multiprotocol network can be a complex task. HP routers simplify this task with a Quick Configuration utility. Quick Configuration allows you to configure the minimum required parameters to enable links and routing/bridging services in the most common configurations. Parameters are entered or modified in a spreadsheet-style configuration summary screen, with the aid of context-sensitive help and error information.

Configuration Editor Network links in a basic configuration can often be brought up and verified in minutes. The configuration can be further customized and finely tuned using the Configuration Editor utility. This utility enables the user to modify any configuration parameter required.

More information on the configurators can be found in the user's and reference manuals accompanying your router.

Quick Remote

Nearly every HP router includes an easy-to-use utility for automatically configuring the remote routers attached to its WAN ports. This is called Quick Remote. It is a simple screen that guides the network administrator at a central site to input and store the parameters minimally required to configure bridging, IP routing, and IPX routing on the HP routers attached on the point-to-point WAN links. Then these remote routers can be unpacked, installed, cabled, and switched on (by personnel with no special networking expertise). Each router immediately sends out a Bootp request and automatically acquires the configuration from the central router on its WAN link (where Quick Remote was used). The remote router then uses the supplied configuration to begin routing with no operator intervention. For each remote router, Quick Remote sets up the WAN link and the first LAN port.

(Even if IP and IPX routing services are not needed, Quick Remote aids in configuration of the other services that are needed. For example, if only AppleTalk routing is being used, the central site administrator could minimally configure IP and Telnet services. After the remote router begins routing, then the administrator could make a Telnet network connection to the remote router, and configure AppleTalk and boot the router using Quick Configuration.)

Note

The HP Router PR does not have Quick Remote. Quick Remote can be used to supply the configuration to any HP router except the HP Router 650. Some types of WAN links (such as X.25, frame relay, etc.) require more configuration than the default point-to-point configuration supported by the automatic startup of branch routers, and thus require more setup than Quick Remote can do at the central site.

Network Management

SNMP Management

HP routers have a Simple Network Management Protocol (SNMP) agent, enabling them to be managed locally, using a console, or remotely, using a modem or Telnet (remote terminal access to console commands) or network management applications. The routers support most standard MIB-I variables, as defined in Request For Comments (RFC) 1156, as well as HP-specific MIB variables. The SNMP agent and the MIB are used when viewing routing and address translation tables or statistics with a console session.

All HP network management products can autodiscover the router and access the router's MIB variables. Additionally, the HP OpenView Interconnect Manager/UX product has router-specific management applications to monitor and configure the HP routers.

HP EASE Instrumentation

HP routers have an EASE (Embedded Advanced Sampling Environment) agent that makes the router a traffic sampler for the connected LANs and WANs. An in-depth, accurate representation of network traffic can be sent, using SNMP, to a network server—with less than ¼ of 1% network overhead. Then EASE applications can be used, standalone or integrated into HP OpenView, to analyze this information:

- You can constantly monitor six network activities—utilization, frame rate, broadcast rate, multicast rate, CRC and alignment errors—in real time, with customized threshold levels for status reporting. You get advance warning about potential network problems, indicating the problem's location by segment and top nodes. Then you can minimize the impact of top traffic sources by better scheduling or by isolating them with traffic filters or further network segmentation.
- You can store and display traffic patterns over time, with trends of where and what network activity occurs most frequently. Then you can implement additional bandwidth where and when needed.
- You can graph workgroup traffic flows—volume, type, usage, and flow directions—and generate reports of recommended actions to maximize network availability. You can estimate the impact of new applications and services on network capacity.

Architecture and Technology

This product note describes the hardware and software architecture of the Router Series 200, 400, and 600. This note includes the features common to each series; technical data specific to each router is found in the *Release Notes* and *Installation Guide* for each series.

Series 200 and 400 Hardware

Some of the routers included in these series are the following:

Series 200	Series 400
HP Router PR	HP Router ER
HP Router FR	HP Router TR
HP Router TFR	HP Router LR
	HP Router SR
	HP Router BR

The router assembly consists of one or two printed circuit assemblies (PCAs), one or two fans, and a power supply in an enclosure. The enclosure of each of the routers is illustrated in the *Release Notes* for your router.

The PCAs, which use surface-mounted components, include the motherboard, containing most of the hardware components, as described in the architecture section below. Some of the routers contain a smaller daughter PCA as well. The motherboards and the daughter boards each contain some of the port line drivers and connectors. (Ports include LAN, WAN, and console ports.)

The power supply assembly is an “auto-ranging” supply. It will automatically detect and switch for alternating-current (ac) input voltages in one of two ranges: 90–120 volts ac or 200–240 volts ac.

Hardware Architecture

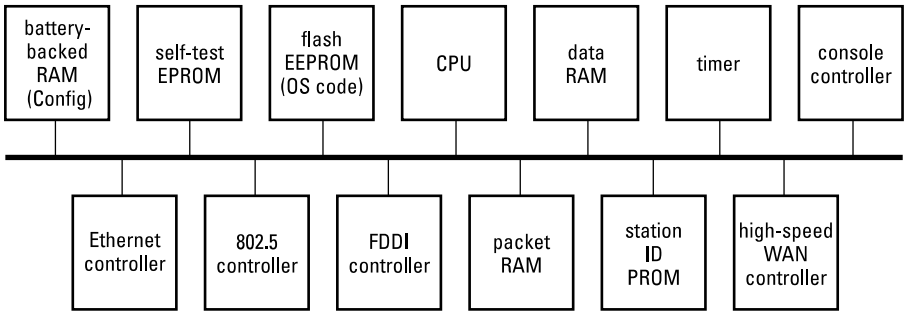

Figure 1. Block Diagram of Series 200/400 Hardware

CPU Most of the series 200 and 400 routers use the Motorola 68020, 68EC020, or 68EC040 processors.

The processor accesses three types of read-only memory (ROM) on the bus:

Flash EEPROM This stores the router operating system code. The use of flash EEPROM technology allows you to update the operating system through the network or the console ports.

Self-Test EPROM This stores the router's self test and console-downloading code.

Station ID PROM This stores the station (MAC) address for each of the router's ports.

The processor accesses two types of random-access memory (RAM) on the bus:

Data RAM Program variables and data critical to packet throughput (such as routing and bridging tables) are stored in a fast memory.

Battery-Backed RAM The router configuration is stored in battery-backed random-access memory.

Timer This generates interrupts and internal clocking for WAN ports.

Console Port Controller

LAN and WAN Controller Coprocessors National's SONIC (Systems-Oriented Network Interface Controller) is a second-generation Ethernet controller that integrates a fully compatible 802.3 encoder and decoder. The token ring coprocessor is Texas Instruments' TMS 380 C16/04, TMS 380 C25, or TMS 380 C26. They support the 4-Mbit/s and 16-Mbit/s data rates as well as universal and local ring addressing. A Motorola FDDI chip set is used for the dual-attached stations on FDDI ports. The MK5025 by SGS Thompson is the WAN coprocessor for serial frame formatting, such as frame delimiting with flags and FCS (frame check sequence) generation and detection. Other information about LAN and WAN ports is found in the "Features of HP Routers" product note in this book, and in the installation guides for your hardware.

Packet RAM Each LAN or WAN coprocessor stores packets received from the line (the network medium) into packet RAM, and transmits packets from packet RAM to the line.

Additionally, there is reset and voltage supervisor circuitry that will reset the router either upon power-up, under low-voltage conditions, or when the Reset button is pressed.

Router 650 Hardware

The HP Router 650 uses a compact table or rack-mountable chassis, illustrated in figure 2. For other detailed illustrations of the router, see its *Installation Guide*.

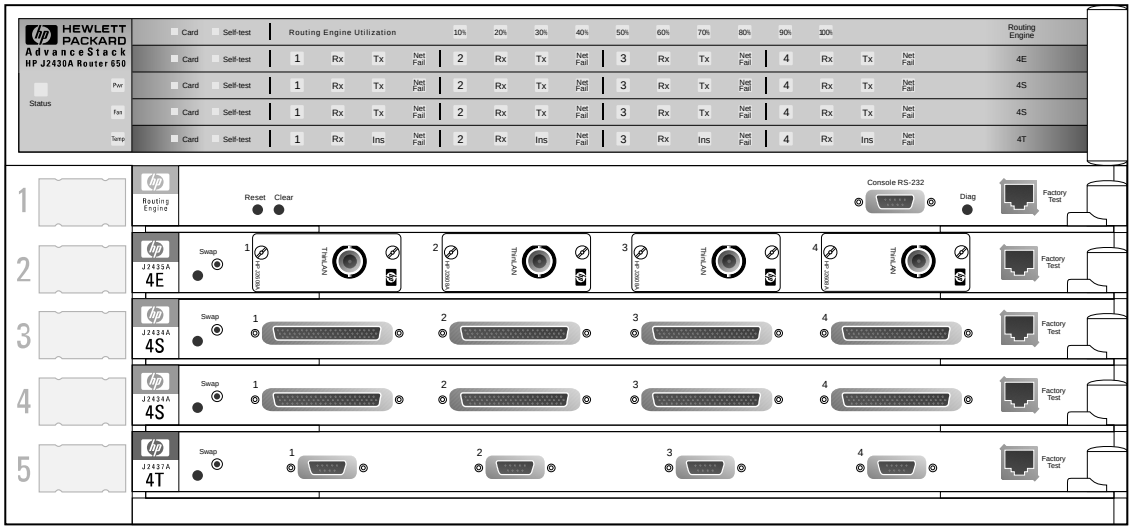

Figure 2. HP Router 650 Chassis

The chassis' first slot is occupied by the system's main processor, here called the routing engine. It includes a console port and router resetting controls. The other four slots may be filled with LAN or WAN interface card modules. Most of these interface products have four network ports each. They can be installed in any combination. The four ports on HP's Ethernet/IEEE 802.3 card product use "garages" to provide flexibility in media attachment. Each garage accepts either a AUI, BNC, FOIRL (fiber-optic), or twisted-pair type of recessed miniature adapter (transceiver) modules. The chassis also contains a module of variable-speed fans and one auto-ranging power supply. A second power supply can be added for load sharing and redundancy.

All interface products (cards and transceivers), the redundant power supply, and the fan module can be replaced on line without interrupting network connections and components that are still operating (here called "hot swap"), and replacements will be autodetected and autoconfigured.

The LEDs are moved out behind all the cabling into plain view. An LED PCA, above the slots, provides multiple status indicators for each port on each card, for overall chassis status, and for the fan and redundant power supply.

Multiprocessor and Memory Architecture

The key feature of the HP Router 650 is its pipelined multiprocessor architecture. The routing engine uses the 33-megahertz Intel i960 CF RISC processor to handle network-layer protocol routing. Each interface card, here called a Data Link Accelerator (DLA) module, also uses a 33-megahertz Intel i960 CF RISC processor to offload the routing engine from the data-link-layer-specific tasks, such as header preprocessing, tabulating data-link-layer counters, token-ring source routing, filtering, and EASE sampling. WAN DLA modules handle PPP, frame relay, X.25, ISDN, SMDS, and compression as well.

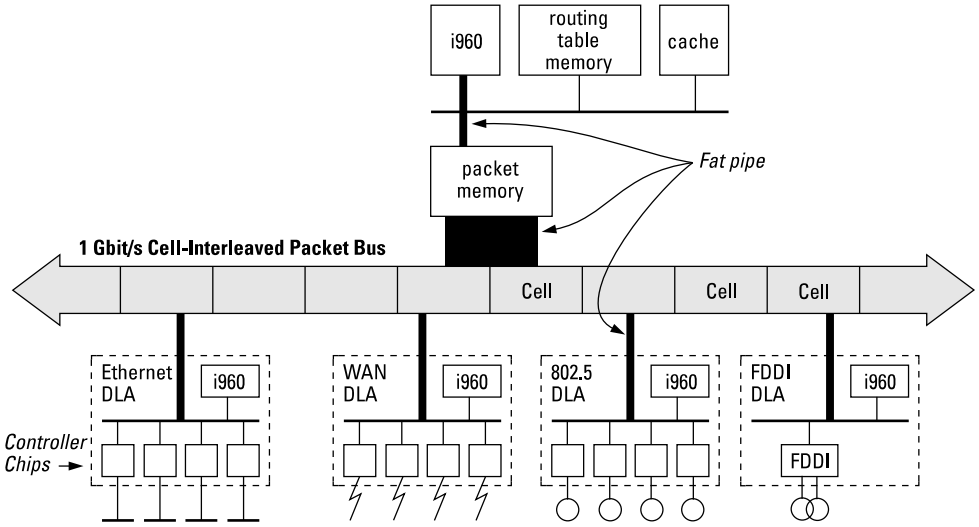

Figure 3. Logical View of Router Series 600 Architecture

The high-speed multiported memory architecture allows the DLA modules to access packets at the same time as the routing engine does, with no degradation in performance. The backplane allows each DLA module to access the multiported packet memory. The memory bus is dedicated solely to switching packets in a cell-interleaved manner, to ensure equitable bus utilitization and to minimize buffering latency.

In addition to the memory bus, the DLA modules also attach to a management bus, used for interprocessor communication and for tasks such as link status notification, “hot swapping” (online replacement of modules), configuration, and all nonswitching tasks.

The routing engine also has 128 kilobytes of high-speed cache to handle even a large number of end-to-end conversation streams, and eight megabytes of routing table memory (expandable by another eight to sixteen) to handle even the largest (1000 or more) router networks. It has a PCMCIA flash card for nonvolatile storage of the preloaded operating system code and of the configuration data. The flash card is removable, and sizes from four to twenty megabytes are supported.

The pipelined multiprocessor architecture is especially powerful when the topology requires routing between various media types; the DLA modules hide the performance-robbing idiosyncrasies of some network media. Ethernet-to-Ethernet routing performance is unaffected by the presence of WAN and token ring source-routed traffic. All WAN-specific processing, such as compression, X.25, and frame relay, is handled by the WAN DLA. The performance of the memory bus is more than one gigabit per second. The throughput of the management bus (whose performance is not critical to operation) is 50 megabits per second. Preliminary testing of overall routing performance demonstrates IP and IPX routing at more than 80,000 packets per second, and bridging at more than 100,000 packets per second.

Data Link Accelerator Architecture

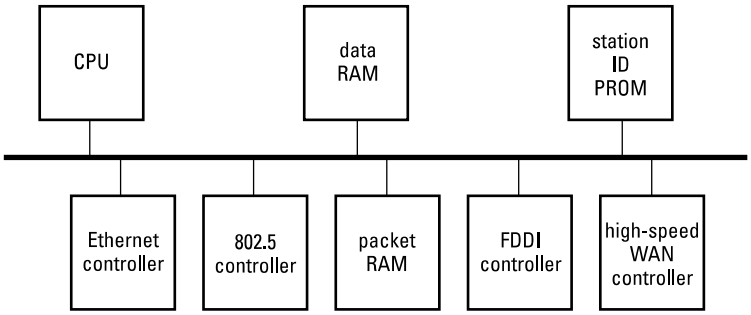

Figure 4. Block Diagram of Data Link Accelerator

CPU The Data Link Accelerators (interface cards) for the HP Router 650 use the 33-megahertz Intel i960 CF RISC processor.

Data RAM Program variables and data critical to packet throughput (such as routing and bridging tables) are stored in a fast memory.

Packet RAM Each DLA coprocessor stores packets received from the line (the network medium) into packet RAM, and transmits packets from packet RAM to the line.

Station ID PROM This stores the station (MAC) address for each of the router's ports.

LAN and WAN Controller Coprocessors National's SONIC (Systems-Oriented Network Interface Controller) is a second-generation Ethernet controller that integrates a fully compatible 802.3 encoder/decoder. The token ring coprocessor is Texas Instruments' TMS380C16/04 or TMS380C25. They support the 4-Mbit/s and 16-Mbit/s data rates as well as universal and local ring addressing. A Motorola FDDI chip set is used for the dual-attached stations on FDDI ports. The MK5025 by SGS Thompson is the WAN coprocessor for serial frame formatting, such as frame delimiting with flags and FCS (frame check sequence) generation and detection. Other information about LAN and WAN ports is found in the "Features of HP Routers" product note in this book, and in the installation guides for your hardware.

Routing Engine Architecture

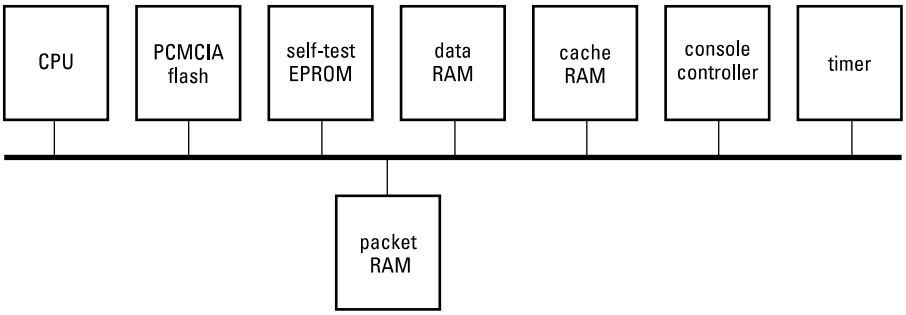

Figure 5. Block Diagram of Routing Engine

CPU The routing engine PCA for the HP Router 650 uses the 33-megahertz Intel i960 CF RISC processor.

Data RAM Program variables and data critical to packet throughput (such as routing and bridging tables) are stored in a fast memory.

Packet RAM Four megabytes of multiported packet RAM are used. Packet sizes up to eight kilobytes can be processed.

Cache RAM 128 kilobytes of high-speed RAM are used for external processor code and data cache.

PCMCIA Flash EEPROM The flash memory on the removable PCMCIA card is used for nonvolatile storage of the operating system code and the configuration data. The PCMCIA card can be replaced to update the operating system code, or to provide more memory for updating through electronic media.

Self-Test EPROM This stores the router's self test and console-downloading code.

Timer This generates interrupts and internal clocking for WAN ports.

Console Port Controller

Routing Software Technology

A router is a layer three—network layer—device. Packets are routed using the network-layer addresses in the network protocol header of the packets. Each “routable” protocol suite, such as TCP/IP or Novell IPX (here called a routing service), manages its own forwarding based on its own address tables, routing protocols, and other routing configuration parameters as a separate software module. Each service's software module is referred to as a redirector.

The HP routers also perform as bridges, which operate at layer two—the data-link layer—independent of and invisible to the network-layer and higher-layer protocols. Packets are forwarded by a bridge using the station addresses (also called physical or MAC or Ethernet addresses) in the data-link-layer header header of the packets. The bridging service manages its forwarding based on address tables and other bridging configuration parameters as a software module of its own. Features for the transparent or source routing bridge, including the spanning tree protocol, are also part of the bridging software module. The redirector for bridging is here called the bridge redirector.

Protocols for the higher layers, such as transmission (TCP) and application (TFTP, FTP) layers, are not used in routing decisions on the router, but support of these protocols depends on the selection of these data-link-layer and network-layer protocols (TFTP requires IP at the network layer), and some of these services are applications provided on the router.

All software modules—redirectors—operate concurrently.

Each physical LAN or WAN link (here called a line)—the cable attached to a port on the router—is also a logical entity called a circuit. A circuit group is another logical entity; it consists of one or more circuits. Each circuit group is the network interface to which each routing or bridging service can forward or route packets. Circuit groups are useful for hiding load sharing and automatic backup from the network layer.

Software Data Flow Architecture

The data flow, that is, the path taken by the packets, for the router software is shown in figure 6. It shows only the software modules critical to router throughput performance—those directly involved with routing or bridging packets through the router. This modular design allows for the addition of new protocols and links. The four basic elements are: driver, data-link service, redirectors, and circuit group manager.

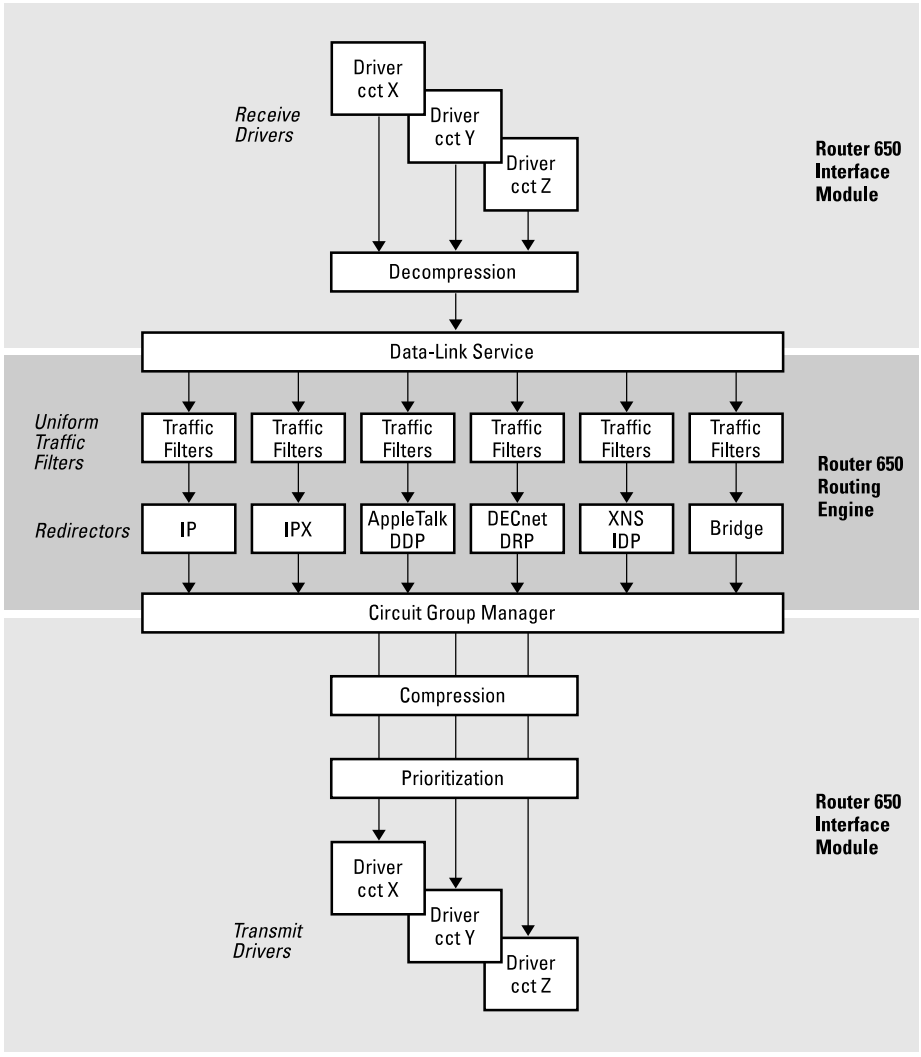

Figure 6. Block Diagram of Software Data Flow (Packets)

Drivers

A driver controls the data-link-layer (layer two) and physical-layer (layer one) protocols. The functions performed by a receiving driver include accepting packets from the network. The functions performed by a transmitting driver include accepting packets from the circuit group manager to forward onto the appropriate LAN or WAN link.

Data-Link Service

The data-link service performs the messaging, resource management, and flow control necessary to move a packet from a source link to a destination link through a routing or bridge module. The functions of the data-link service include demultiplexing incoming packets from the driver, managing internal flow control, and providing transparent access to multiple redirectors concurrently.

Redirectors

Each HP router can support the simultaneous operation of multiple routing and bridging services. Each software module representing one of those services is known as a redirector. Redirectors are responsible for forwarding packets from one circuit group to another. They receive packets from the data-link service. Table lookups result in subsequent modification of the packet header's address and control fields. Redirectors then transmit the modified packets through the data-link service to the specified driver and device(s).

It should be noted that the spanning tree protocol, routing protocols (such as RIP or AppleTalk RTMP), and router management protocols (such as SNMP) are not a part of the performance data path. The routing protocols only collect and maintain the routing information used by the network-layer data-path modules shown in figure 6. The flow of routing information is shown in the next section.

Circuit Group Manager

The circuit group manager manages the circuits and network interfaces (circuit groups). The circuit group manager is responsible for selecting the best circuit from a circuit group using the circuits database. It handles the load sharing and automatic backup functions. The circuits database contains up-to-date metrics about all of the router's circuits.

Software Control-Path Architecture

The flow of routing information used in selection of a route and network interface (circuit group) involves the routing protocols (RIP, ARP, OSPF, etc.) primarily. This flow is different for each bridging and routing service. For some services, the path can involve numerous choices of routing protocol and can involve other features as well, such as the import and export route filters for the IP routing service. The most complex flow is for IP, as shown in figure 7. Figure 8 shows the flow for IPX.

These diagrams of the flow of routing information complement the diagram of the flow of the actual data packets (which are most of the packets transmitted) in figure 6.

IP Routing Information Flow

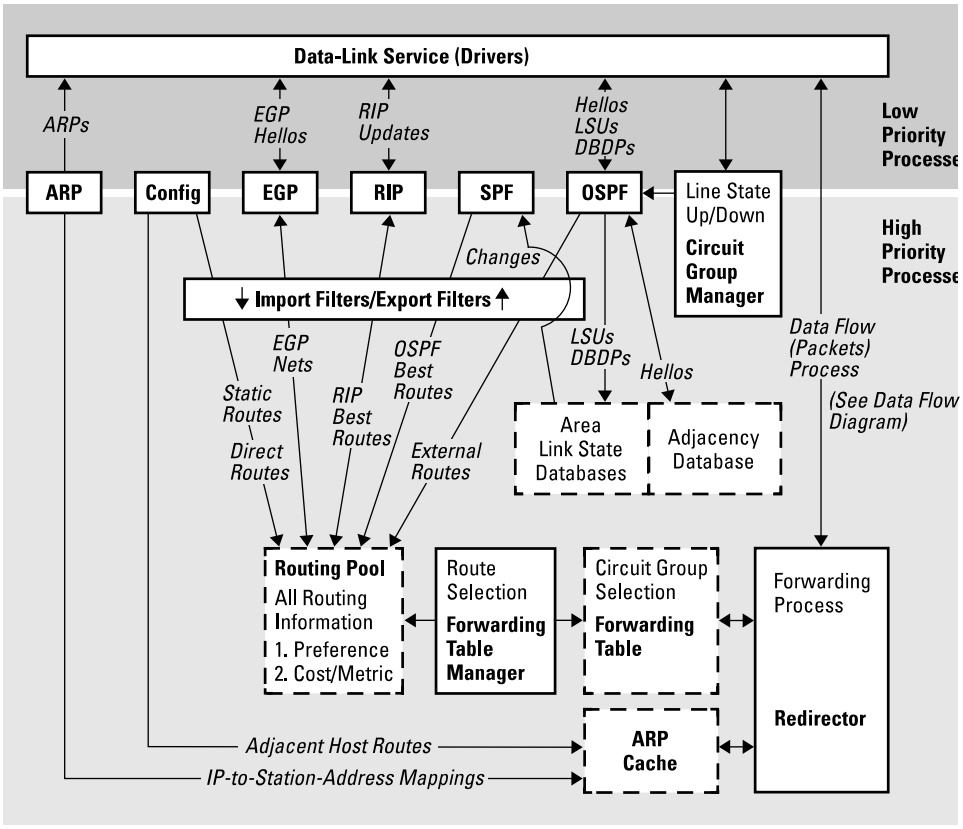

Figure 7. Block Diagram of IP Routing Information Flow

IPX Routing Information Flow

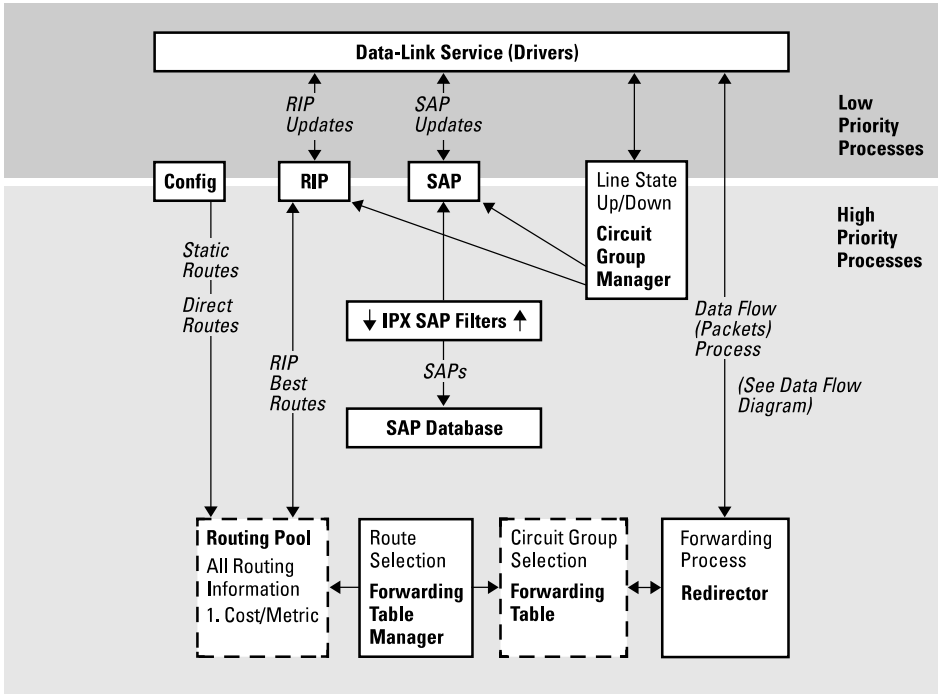

Figure 8. Block Diagram of IPX Routing Information Flow

Branch Office Routing

Introduction

The 1980's were christened recently by the media and vendors alike as the decade of growth for internetworking. The abundance of protocols, cable types, speeds and feeds, as well as unremitting growth, have made designing simple yet comprehensive internetworking solutions the ultimate challenge of the 1990's.

In the coming decade, the industry will be asked routinely to provide customers with multiprotocol support, easy connectivity from home offices to branch-office local-area networks, and uncomplicated management of the remote networks. In addition, managing the costs of installation, use and maintenance of these internetworks will be a prime consideration.

This note covers the issues related to the implementation of branch-office networks, and the steps Hewlett-Packard is taking to help customers build branch-office networks that meet business needs while minimizing costs.

Overview

PC and workstation LANs are proliferating in branch offices for many reasons:

- to take advantage of low-cost PCs and new client-server applications;
- to share expensive peripheral devices;
- to customize applications to meet local needs; and
- to accommodate workgroup computing styles.

This trend is forcing corporate network planners to consider solutions for linking the often large number of branch offices to the corporate information network. The problem is how to do it reliably, yet cost effectively. Compounding the problem is the fact that there are typically no MIS resources at the branch office.

Just as every corporation's networking requirements are different, so are the solutions for each environment. Connecting branch offices, for example, may be as simple as establishing a single Ethernet link per remote office back to a corporate or regional site using a leased line, as shown in figure 1.

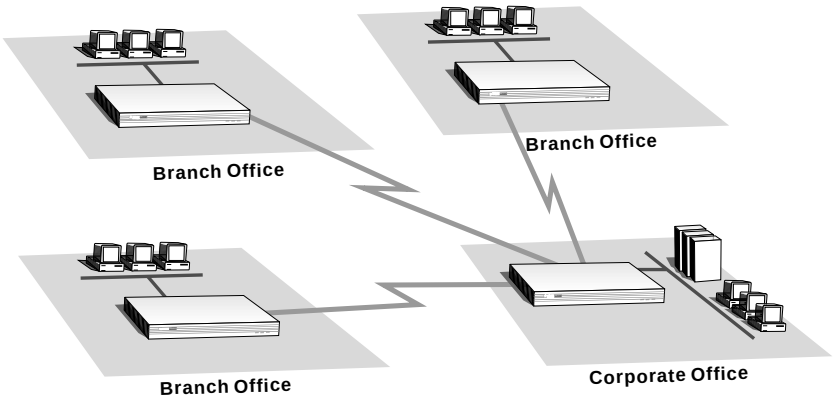

Figure 1. Branch Office Network

Alternatively, many businesses, such as banking, financial services and insurance, require more than a single LAN at the remote office, as shown in figure 2. And there may be legacy devices there, such as IBM terminal controllers or automated teller machines using X.25 or SDLC, performing mission-critical functions, that must also be included in a particular branch networking solution. For additional reliability in the event of link failures, dial backup may be a consideration.

Clearly, branch office communication requirements vary. What doesn't vary is the need for a cost-effective solution that meets the networking needs of the business, and that is as simple as possible to administer.

Branch Routing Requirements

To develop cost-effective solutions for branch office networks, network designers must determine the basic communications requirements of the applications for which the network is being designed, including the underlying protocols and the relative priorities of applications based on mission criticality. Next, but more importantly, designers must understand the factors that affect the long-term cost of ownership of a network to keep costs to a minimum. HP's strategy is to minimize the costs associated with implementing and managing branch office networks. The major areas of focus include:

- Controlling WAN Costs
- Controlling Administration Costs
- Controlling and Protecting Equipment Investments

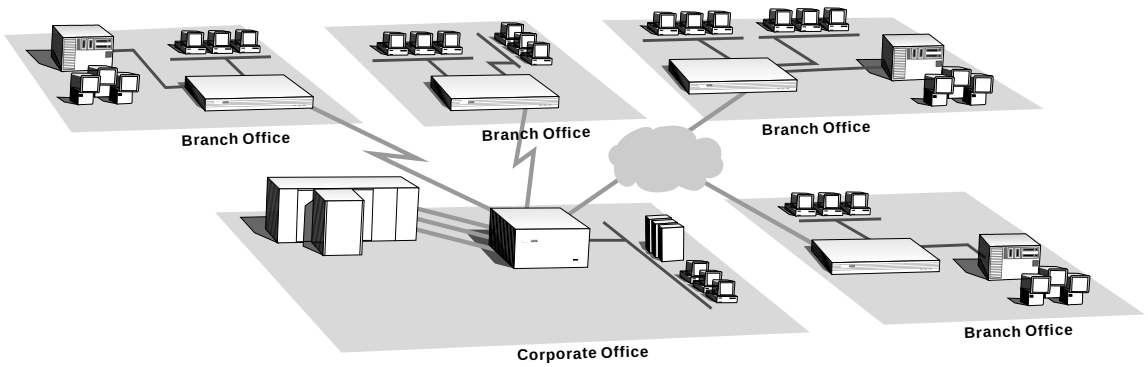

Figure 2. High Complexity Branch Office Network

Controlling WAN Costs

WAN communication charges are now and will continue to be the single greatest cost component in annual branch office networking budgets. Support for popular WAN protocols and features that improve WAN performance, add robustness, and eliminate the need for redundant links all play a significant role in reducing WAN communications costs.

Improving Performance

HP has several new features that improve WAN performance, add robustness, and get the most out of wide area communications facilities, avoiding additional communications charges wherever possible. These include:

- HP packet-by-packet compression
- Prioritization/bandwidth reservation
- Latency control

Packet-by-Packet Compression

Packet-by-packet compression (HP PPC) is an innovative HP data-link compression technology that improves throughput and can eliminate the need to add costly incremental bandwidth. Unlike other compression solutions, HP PPC was developed for use with multiprotocol data over all link types including leased lines, packet switching networks, and circuit switching networks. Because other compression solutions require a reliable link protocol such as LAP-B, they are not suitable for use on links, such as frame relay, that use an *unreliable* link-level protocol. HP PPC does not require special protocols and will thus be available for use on any link type.

With HP PPC, compression is achieved using very small dictionaries that conserve memory. This helps make compression achievable (cost-effective), even on routers with many synchronous ports.

Additionally, HP PPC protects against data expansion. If the resulting compressed packet is larger than the uncompressed packet, the packet is sent uncompressed. A typical example is when the packet being transmitted is already compressed.

Like other compression schemes based on the Lempel-Ziv algorithm, the actual compression ratio obtained depends on the type of data being transmitted. Generally, the compression ratio obtained over a long period of time averages about 1.8 to 1.

HP PPC is implemented in software and is currently supported over leased lines at link speeds up to 64 Kbit/s. Compression will be available on dial-up circuits and frame relay circuits in 1994, and on X.25 in 1995.

Prioritization/Bandwidth Reservation

One of the characteristics of extended LAN networks is that traffic is inherently *bursty*. This means that varying amounts of data are transmitted at irregular time intervals. This can have adverse effects on certain types of applications–especially those that rely on bridging of connection-oriented protocols such as IBM's SDLC. Problems such as connection timeouts and unpredictable response times may occur when connectionoriented data is queued waiting for the transmission of other data.

HP's traffic prioritization is a feature that a network designer can use to ensure that delay-sensitive data is given priority over data that is less sensitive to delay. For example, SDLC traffic that is bridged from a terminal controller to a front-end processor can be given priority over other routed traffic.

Bandwidth reservation works with traffic prioritization. Protocols or data can be prioritized, and during periods of peak WAN utilization, bandwidth can be reserved for each priority level. The goal is to ensure that the highest priority data receives enough bandwidth without *starving* applications transmitting lower priority data.

Controlling Latency

Latency is the time a packet spends inside a router as it is forwarded and queued for transmission. In router-based networks, the *apparent time* required for an application process to execute across a network is often determined by link propagation delay and router latency, latency being the primary factor.

Latency can be reduced by tailoring the size of the transmit buffer queue of a WAN link to fit the speed of the link. HP routers improve application responsiveness by providing network designers with precise control over the amount of data that is buffered on WAN links, thereby reducing latency.

Reducing Cost and Improving Robustness

Once relegated to the communications backwaters, circuit switching (dial-up) technology is becoming an increasingly attractive proposition for branch office networking. The near universal availability of switched digital communications (Switched 56 and/or ISDN) as well as the new CCITT V.fast recommendation clearly put circuit switching in the mainstream as far as bandwidth is concerned. Circuit switching can be an ideal solution both as a primary and backup connection mechanism. HP has developed several unique solutions using dial-up communications to help customers solve branch office networking needs.

Dial on Demand

HP routers can use dial-up communications as the primary means by which data is exchanged between branch offices and corporate or regional offices. While not the ideal communication facility for transaction-oriented applications, dial-up is very cost effective for batch networking requirements such as file transfer and electronic mail. HP routers can use both V.25 bis and manual (DTR mode) modems, and terminal adapters.

Call Controls

Dial-up connections are governed by “smart” connection controls to take maximum advantage of dial-up circuits. Call controls allow a network designer to configure a router with billing period information. This information is used to hold dial-up circuits open as long as possible when it can be done at no cost, for example, until the end of a three-minute initial billing period. This often avoids a second call and additional dial-up charges.

Dial Backup

One of the traditional strengths of routing technology is the ability to provide robustness using alternate routes. Dial-up links are a very cost-effective way of providing alternate path capability at the branch office level. HP has several dial backup solutions. For multiprotocol applications over leased lines, HP has a backup circuit capability wherein the backup circuit is only activated when its primary circuit counterpart has failed.

For IP routing applications, HP has a unique backup solution that uses a small number of backup circuits (and modems) at a regional or corporate site to service a much larger number of remote sites with dial backup capability. This allows a network designer to provide only as many backup ports as there are likely to be primary circuit failures at any one point in time.

Planning & Controlling WAN Capacity

Understanding wide area networking traffic volumes and characteristics is critical to controlling WAN costs. If a 64 Kbit/s link can adequately satisfy user needs instead of a T1 or E1 link, WAN line, equipment, and administration costs can be dramatically reduced without affecting service levels.

Few solutions exist to understand WAN capacity utilization and the ones that are available are expensive, place significant overhead on the network, and are difficult to set up and adminis-ter in a remote branch environment except on a very limited basis. What is needed is for this capability to be transparently integrated into the network devices themselves. HP routers provide just this capability.

All HP routers include a network traffic sampling capability called HP EASE. EASE stands for Embedded Advanced Sampling Environment. This sampling technology provides an in-depth, accurate representation of network traffic without additional special equipment or set up and with virtually no additional network overhead (less than 1/4 of 1%).

Powerful graphical applications provide the tools to analyze information on network traffic. They are available as standalone applications or integrated with the HP OpenView environment. With these applications, the network administrator can:

- Better manage exisiting WAN capacity—EASE applications help the administrator identify top sources of network traffic by network node, communicating pair, protocol, services or application. With this information, the network administrator can work to minimize the impact of top traffic sources by isolating them (i.e. using traffic filters or with further network segmentation) or through better scheduling of network activity thus avoiding the need to add costly WAN bandwidth.
- Understand network traffic trends to anticipate future capacity requirements—This allows the network administrator to carefully manage the implementation of additional capacity to ensure bandwidth is added only when needed and in a manner that guarantees adequate service levels to end users.
- Estimate the impact of new applications and services on WAN capacity—As new applications and services are implemented, the network administrator can accurately gage the bandwidth requirements of these new capabilities to ensure a smooth rollout with the least cost.

These applications also provide information on network health to facilitate fast isolation and resolution of network problems. The applications capture traffic trends in the network and warn of impending problems. This helps ensure continuous network operation.

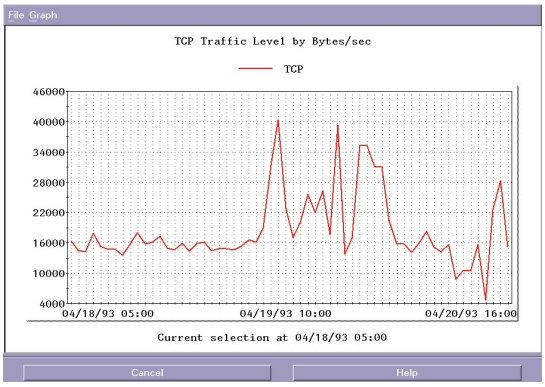

Figure 3. Traffic volumes between remote sites over a specified time period.

Controlling Administrative Costs

Extending the corporate network to remote offices concerns many network managers who are worried about the supportability of a branch office network. Branch offices do not have MIS network administration staff on site to assist in the implementation or ongoing management of the network. Under these conditions, routers must be installable and manageable by MIS administrators remotely. Furthermore, ease-of-use features that simplify installation and management are especially important because of the cumulative effect they have on reducing the MIS administrative burden. Simplifying the administration of router-based networks is an area in which HP is making significant contributions and will continue doing so.

"Instant-On" Branch Office Router Installation

Imagine the ideal solution for making router products installable in remote branch offices. The solution would merely require the router to be placed on a shelf, communication cables connected, and the power to be applied. To be failsafe, all else should be automatic. There should be no tools required, no disks to insert, no switches to set, no terminals to connect, and no software to configure. This ideal solution is available from HP today!

Hardware Designed for Branch Offices

Every HP router, from the HP Router ER to the recently announced HP Router 650, have the following features to enhance remote support:

- Software preloaded in flash EEPROM memory. No disks to get lost, misplaced, or fail. The software is always loaded, and it can be updated over the network.
- Absolutely *no* hardware configuration options. Hardware can *never* be misconfig-ured. *All* configuration options can be controlled remotely through software.
- Automatic interface detection circuitry. HP routers automatically detect the type of WAN interface being used. To upgrade an interface from RS-232 to V.35, merely replace the cable; all HP routers have circuitry to support all of the commonly used physical interfaces, and HP routers automatically detect what type of interface is being used. Again, *no* switches to set, *no* configuration to change.

Software Designed for Automatic Branch Office Installation

Dynamic link configuration and SmartBoot help speed the installation of branch office routers and eliminate costly and time consuming configuration errors.

Dynamic Link Configuration

In the fall of 1993, HP introduced dynamic link configuration for wide area interfaces using leased lines. With dynamic link configuration, leased-line interfaces no longer require the configuration of data-link-layer parameters. When two routers are connected to an operation-al link, the routers automatically negotiate all link parameters and establish a connection. While the connection is being established, the link speed is automatically determined. Transmit queues are automatically configured based on the link speed to minimize latency and improve application responsiveness.

Dynamic link configuration works not only on router-to-router links but also on router-to-remote bridge links. This is especially useful when connecting HP Remote Bridges in branch offices to a router such as the HP Router 650 in a regional or corporate office.

SmartBoot

A companion feature introduced in the first half of 1994, SmartBoot, allows a branch office access router to automatically retrieve its routing configuration information after establishing a connection with a corporate or regional router. With SmartBoot, routers at branch offices can literally be installed in minutes by people with no special training.

With SmartBoot, routers at branch offices can literally be installed in minutes by people with no special training.

SmartBoot is operating code that is implemented in both corporate/regional and branch office routers. Quick Remote, one of the components of SmartBoot, lets a network administrator create and store configuration files on a central site router for its attached branch office routers. All HP routers with the exception of the HP Router PR can store one remote configuration per WAN interface. The configuration files created with Quick Remote include high-level information such as addresses and which of the bridging and routing (IP and/or IPX) modules are to be enabled. Two interfaces (one LAN, and one WAN) can be configured per remote router using Quick Remote. SmartBoot is ideal for routing applications using routers like the HP Router FR and PR at branch office locations.

Branch office routers acquire their configuration when they are placed in service. When the router first boots up, it sends a Bootp request on its WAN interface. If the attached corporate router has a stored configuration for the branch router, it is sent when the Bootp request is received. The branch office router uses this configuration to complete its bootup sequence, after which the branch office network is operational.

Controlling and Protecting Equipment Investments

Equipment cost is a major concern in the deploy-ment of branch office networks that may contain tens or hundreds of routers. Seemingly small differences in purchase price can represent large savings as networks grow. Warranty is another important source for reducing costs. Longer warranty periods delay the need for hardware support agreementsan important source of cost savings as the network is starting up.

A Broad Range of Cost-Effective Solutions

From the introduction of the Router ER (HP's first router), HP has led the initiative to reduce the cost of routing platforms. Branch solutions such as the HP Router PR start at $2395 ranging to the newly announced high-end HP Router 650 with an entry level configuration priced at $12,000.

The HP Router 650 extends HP's family of routers to address the specific requirements of a headquarters central office. Designed to offer the highest performance in its price class, this new system delivers performance comparable to other high-end products that cost two to three times more.

Figure 4. HP Router 650

Warranty

HP provides a full 3-year warranty on all router products–the longest warranty in the industry. Besides lowering equipment costs, this is also a statement about the quality of HP's router products and their ability to perform year after year.

Future Directions

HP is working to make branch office routing even more cost effective, and easier to implement and manage. Capabilities expanded to draw more performance out of the WAN will include:

- Additional compression algorithms optimized for specific link types including packet switched networks.
- More traffic control and management features to improve application access to WAN resources.
- New WAN traffic planning and analysis tools.

To help control and contain the administrative costs associated with implementing and managing branch networks, HP is developing:

- Instant-On installation for routers in packet switched environments.
- New management tools that will both speed the design and implementation of router-based networks as well as simplify on-going management. New features will include graphical network design, configuration validation, and centralized configuration and software distribution.

2

Routing Services Notes

Routing Services Notes

- Bridging Service
- Internet Protocol Routing Service
- Novell IPX Routing Service
- AppleTalk Phase 2 Routing Service
- DECnet Routing Service
- A Primer on HP Probe
- Data Compression for WAN Links

Bridging Service

The HP router can operate as a multiport bridge. The bridging service may be enabled independently of any of the routing services; if not enabled, then the router will discard packets with protocol types not enabled or supported (such as IBM SNA or DEC LAT). It is possible to use the HP router solely as a bridge by not enabling any of the routing services.

This bridging service operates at the media access control (MAC) sublayer of the data-link layer of the OSI reference model.

As an Ethernet/IEEE 802.3 transparent bridge (or learning bridge), an HP router supports a variety of filtering options at the MAC layer, such as source and destination address and protocol type. The IEEE Spanning Tree Protocol (STP) is available for management of bridged networks with mesh topologies. STP allows a bridged network to have redundant paths. In the event of a primary link failure, a backup link takes over, thereby ensuring continued data transmission between all reachable network segments.

Translational bridging is used between FDDI and either Ethernet or IEEE 802.3 networks. Source-route translational bridging is used for bridging between token ring (IEEE 802.5) and Ethernet or IEEE 802.3 networks. These addressing schemes differ at the data-link layer, specifically at the MAC layer.

Source-routing bridging is used to connect token rings that contain systems communicating with non-routable protocols such as IBM SNA or NetBIOS. Thus, when an HP router is configured for source-routing bridging, it can be used in any application that would otherwise be performed by an IBM Source Route Bridge. Transparent (also known as learning) bridging is always enabled with the bridging software, so that when source-routing bridging is performed concurrently, the HP router is known as a source-routing/transparent (SRT) bridge.

Transparent Bridging

Transparent bridges provide network interconnection and/or extension services to LANs that employ identical protocols at the data link and physical layers. Transparent bridges place no burden on end nodes; they take no part in the route discovery or selection process. From the point of view of an end node, it appears that all nodes are resident on a single extended network with each node identified by a unique MAC-level address. Essentially, a transparent bridge provides a relatively uncomplicated relay function.

The transparent bridge provides three primary services: it learns the addresses of end nodes on connected networks; it relays frames on the basis of its acquired knowledge of end-node addresses; and (if the spanning tree algorithm is enabled), it ensures a loop-free topology throughout the extended network.

The transparent bridge identifies nodes on a network using the station address (also known as the MAC address, Ethernet address, and physical address). The end-node's station addresses is found in the source address of each frame received by the bridging service. As the bridge receives frames, it builds and updates a database that lists each source address. This bridging database is known as a forwarding and filtering table or an address table; it is comparable to the "routing table" used for the other services that actually route at the network layer. In the address table, each source address is accompanied by the circuit group on which the address was observed and by a timer value that indicates the age of the observation.

The transparent bridge relays frames on the basis of address table entries. When the bridge receives a frame, it compares the frame's destination address with addresses found in the address table. If the bridge fails to find a match, it relays the frame on all circuit groups (except the circuit group on which the frame was received). This action of relaying a frame on multiple circuit groups is called flooding.

If the bridge finds a match between the destination address and a address table entry, it compares the circuit group on which the frame was received with the circuit group associated with the table entry. Identical circuit groups indicate that the source and destination end nodes are located on the same physical network. In this instance, because relay is not necessary, the bridge drops the frame. Different circuit groups indicate that the source and destination end nodes are not located on the same physical network. In this instance, the bridge relays the frame on the circuit group found in the address table.

With the spanning tree algorithm enabled, the transparent bridge ensures a loop-free topology by cooperating with other bridges in the extended network. The algorithm provides a single path (composed of bridges and intervening LANs) between any two end nodes.

Spanning Tree Algorithm

The IEEE 802.1 committee has issued a standard applicable to all MAC-level bridges. Much of this standard is concerned with the operation of bridges in topologically complex environments which may contain parallel (also called redundant) bridge connections between multiple LANs. Such parallel connections cannot be tolerated within a transparent bridging environment.

For example, in figure 1, the red and white LANs are connected by two routers serving as parallel bridges, bridge 1 and bridge 2. Consider the chain of events when end node J on the red LAN first sends a frame to end node K on the white LAN. The frame originated by end node J and addressed to end node K is read by both bridge 1 and bridge 2. As this is the first frame between J and K, the address table of neither bridge contains an entry for J or K.

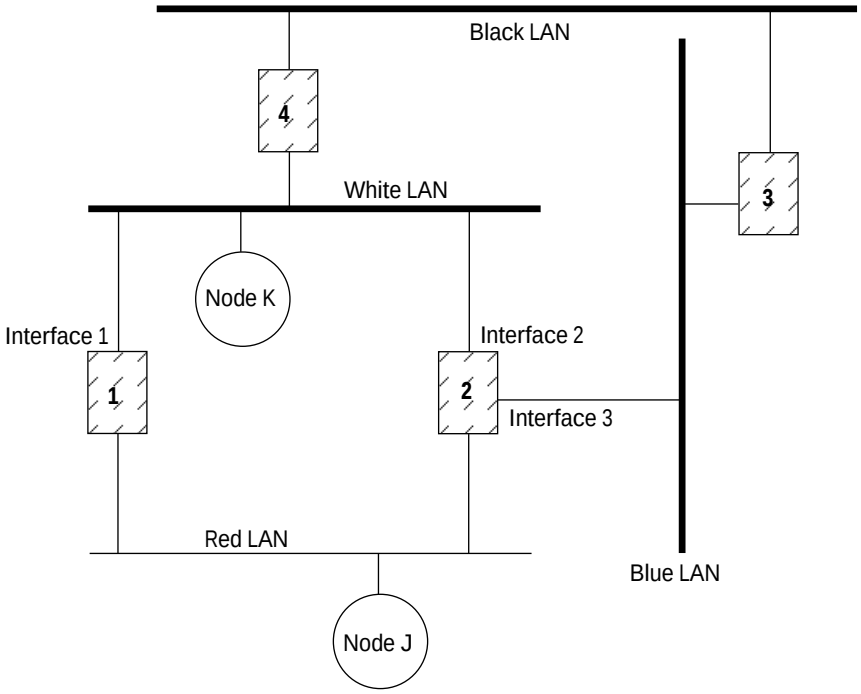

Figure 1. Parallel Bridge Topology

Each bridge updates its address table to indicate that J is in the direction of the red LAN. After updating its address table, each bridge floods the frame: bridge 1 relays the frame over interface 1 and bridge 2 relays the frame over interface 2. Bridge 2 also relays the frame over interface 3; to simplify the example, however, this frame will not be traced.

Next, end node K receives two copies of the frame originated by end node J. While the reception of duplicate frames by a node is not generally fatal, at best such duplication represents an inefficient use of available bandwidth. Of graver consequence is the effect of duplicate frames on bridge 1 and bridge 2. The frame flooded by bridge 1 onto interface 1 is ultimately read by bridge 2 on interface 2. When bridge 2 reads this frame, it updates its address table to indicate end node J is in the direction of the white LAN. In a similar fashion, bridge 1 reads the frame flooded by bridge 2, and it updates its address table to show end node J in the direction of the white LAN. Consequently, the address tables of both bridges are corrupted and neither bridge is now able to properly forward a frame to end node J.

This corruption is caused by the existence of alternate routes between hosts. Such alternate routes are generally referred to as loops. The spanning tree algorithm (fully described in IEEE 802.1 MAC Bridges) ensures the existence of a loop-free topology in networks that contain parallel bridges. The algorithm provides a single path (composed of bridges and intervening LANs) between any two nodes in such an extended network. It also provides a high degree of fault tolerance by allowing for the automatic reconfiguration of the spanning tree topology in the face of bridge or data-path failure. Five management-assigned values are required for derivation of the spanning tree topology:

- A multicast address specifying all bridges within the extended network
- A network-unique identifier for each bridge within the extended network
- A unique identifier for each bridge/LAN interface (called a port)
- A priority specifying the relative priority of each port
- A cost for each port

With these values assigned, bridges broadcast and process formatted frames (called bridge protocol data units or BPDUs) to derive a single loop-free topology throughout the extended network. BPDU frame exchange is accomplished quickly, thus minimizing the time du which service is unavailable between hosts.

In constructing a loop-free topology, the bridges within the extended network first determine the root bridge, the bridge with the best (that is, lowest) priority value. This bridge serves as the root of the loop-free topology.

After determining the identity of the root bridge, all other bridges calculate path costs, that is the cost of the path to the root bridge offered by each bridge port. Each bridge designates the port that offers the lowest-cost path to the root bridge as the root port. In the event of equal path costs, the bridge designates the port with the best (that is, lowest) priority value as the root port.

On each LAN within the extended network, one bridge (the one whose root port offers the lowest-cost path to the root bridge) is selected as the designated bridge. The port that connects the LAN to the designated bridge is selected as the designated port. This port—said to be in the forwarding state—carries all extended network traffic to and from the LAN.

This process ensures that all redundant ports (those providing parallel connections) are removed from service (placed in the blocking state). In the event of a topological change, or in the event of bridge or data-path failure, however, the algorithm derives a new spanning tree that may move some such ports from the blocking to the forwarding state.

Using figure 1 as an example, the implementation of the spanning tree algorithm could remove bridge 1 from service and block bridge 2/interface 3. Figure 2, below, shows the resulting logical topology that provides a loop-free topology with only a single path between any two end nodes.

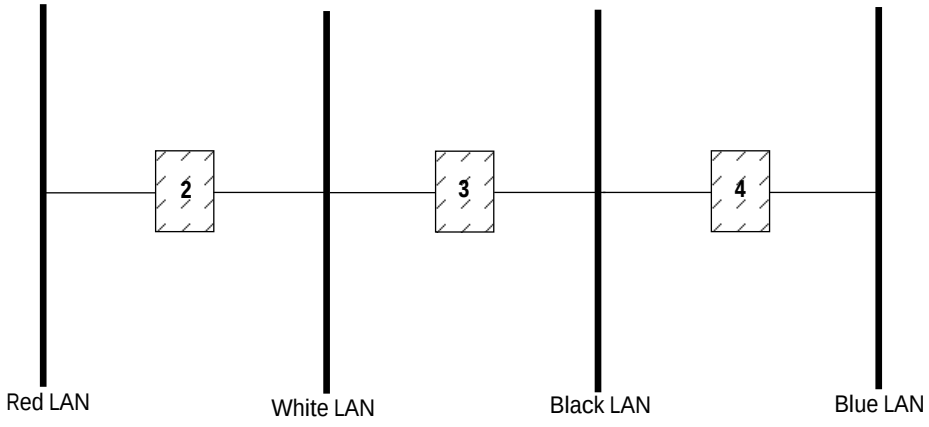

Figure 2. Spanning Tree (Loop-Free) Logical Topology

Token Ring Solutions

HP router software provides support for IBM source-routing bridging. This makes many new routing and bridging solutions possible in token ring environments. This section examines these routing and bridging applications and explains concepts central to routing and bridging in a token ring environment. A few token ring limitations are described at the end of the section.

The Token Ring Network

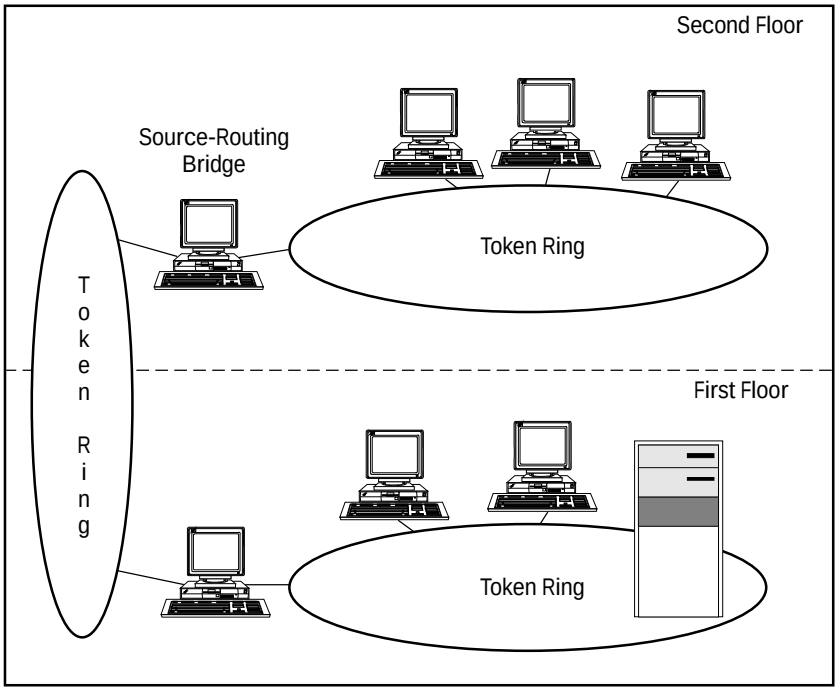

Figure 3. Typical IBM Token Ring Network

Figure 3 shows a typical IBM token ring network.

The backbone ring (the building's backbone network) is a 16-Mbit/s token ring, which is accessible from each floor. Application rings, installed on each floor, are 4-Mbit/s token rings. Computers are attached to the application rings. IBM source-routing bridges are used to connect the application rings to the backbone ring. Source-routing bridges are PS/2 personal computers with two token ring interface cards. The functions of the source-routing bridge are provided by IBM's Token Ring Bridge Program.

Source Routing

Source routing is just what the name implies. In a multi-ring topology where token rings are connected by source-routing bridges, originating systems (end systems or nodes) specify the route that packets must follow to reach their destination. To determine the route to a particular destination system, a system must first discover the route to that destination. An end system with data to send to another end system (with which it has not yet communicated) begins the discovery process by transmitting a discovery packet. All source-routing bridges receive the packet, insert their network identifiers in the packet's Routing Information Field, and then retransmit the packet on all network interfaces other than the one on which the discovery packet was received. The destination system receiving the packet responds by returning the route information received in the discovery packet to the originating system. Once the route has been discovered, originating systems send data to the destination with the route inserted in the packet's Routing Information Field.

Unlike transparent (or learning) bridges, such as the HP 10:10 LAN Bridge, source-routing bridges do not maintain an address table. Thus, bridging decisions are not based on a forwarding table. Rather, forwarding and filtering decisions are based solely upon the Routing Information Field contained in each packet. The burden of discovering and defining the route a packet will take rests mainly with the communicating systems. Source-routing bridges assist in the discovery process and subsequently follow routing instructions contained in each packet.

In contrast, systems using transparent bridges have no information indicating whether a destination system is on the LAN or not. The burden of discovering a route and routing packets to a particular system rests entirely with the transparent bridge.

Routing in Token Ring and Mixed-Media Environments

The IP, IPX, XNS, and AppleTalk routing services provide support for networks containing source-routing bridges. Source-routing support has not been provided for DECnet. Therefore, the following discussion of routing applies only to IP, IPX, XNS, and AppleTalk.

Source routing may be enabled or disabled for each routing service on each interface. The routing services act as end systems (the PC or the mainframe in figure 3), from a source-routing perspective. This is what is meant by support for source-routing bridges. The routing services never function as source-routing bridges.

The routing services are users of the underlying physical and data-link layers. Thus, the routing services do not modify data-link-layer headers (source and destination station addresses, source-routing information, etc.).

Since the routing services' operations are independent of the data-link layer, there are no media-dependent, topological restrictions on the use of routers in large mixed-media networks. The network topology shown in figure 4 illustrates this point. Each computer in the network is able to communicate with any other computer. This assumes that all systems communicate using routable protocols supported on token rings (IP, IPX, XNS, and AppleTalk), and that the routers are configured to route (not bridge) these protocols.

On HP routers, the bridging service as well as the routing services listed above all support networks with token rings and mixed media. Routing is generally better at handling them than bridging, so give preference to routing if you can. However, the following two conditions will require you to use bridging instead of routing.

- Non-routable protocols (such as SNA, NetBIOS, 3270) are used on the network.
- The network is not sermented properly to allow routing.

Most of the rest of this bridging note describes how the bridging service handles source routing and mixed media.

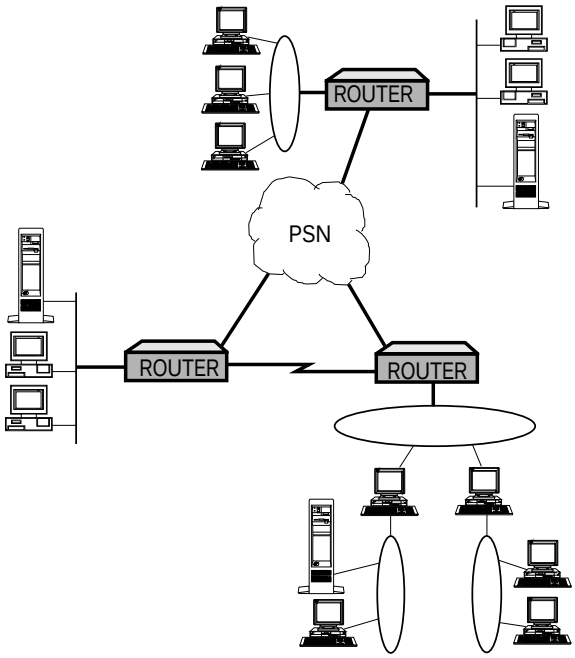

Figure 4. Routing in a mixed-media environment permits communication between all systems.

Bridging in Token Ring Environments

The bridging service on the HP routers includes source-routing bridging. Source-routing bridging is used to connect token rings containing systems that communicate using non-routable protocols such as IBM 3270 or NetBIOS. Thus, the HP router, when configured for source-routing bridging, can be used in any application that would otherwise be performed by an IBM Source-Routing Bridge.

Source-Routing and Transparent Bridging

When bridging software is enabled, transparent bridging is always enabled. Optionally, the bridge will also perform source-routing bridging. Bridges that perform both types of bridging concurrently are referred to as source-routing/transparent (SRT) bridges. Thus, when source-routing bridging is enabled, an HP router can function as an SRT bridge.

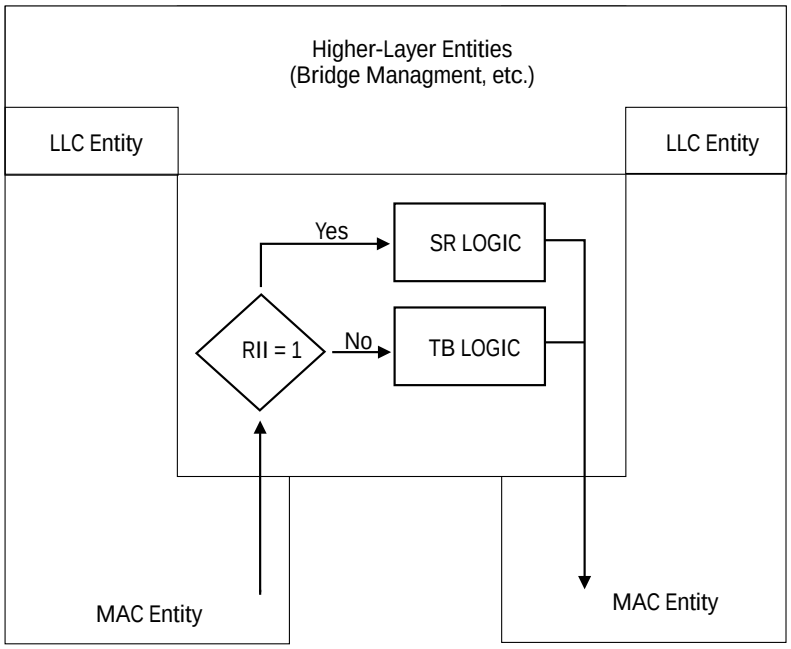

Figure 5. Architecture of the Source-Routing/Transparent (SRT) Bridge

Figure 5 shows an architectural model of the SRT bridge. When the SRT bridge receives a frame, it examines the Routing Information Indicator (RII). Frames with RII=1 are processed using source routing (SR) logic. Frames with RII=0 are processed using transparent bridge (TB) logic. The RII is the first bit transmitted of the source station address (MAC address). This bit occupies the same position in the source address as does the individual/group bit in the destination station address.

Source-Routing/Transparent Bridge Versus Source-Routing Bridge

A source-routing (only) bridge provides inter-ring communication only for systems that support source routing. Systems that do not support source routing can communicate only with systems on the same token ring. An SRT bridge, in contrast, provides inter-ring communication for systems that support source routing as well as those that do not support source routing.

Limitations An SRT bridge does not, however, provide communication between a source-routing system on one ring and a non-source-routing system on another ring.

There are some basic differences at the MAC layer between token ring/IEEE 802.5 and Ethernet. As a result of these differences, the SRT bridge does not support bridging between systems on dissimilar LANs. Traffic between systems located on token ring LANs, therefore, cannot be bridged to systems on Ethernet/IEEE 802.3 LANs.

Bridging in Mixed-Media Networks

Differences Between Token Ring and Ethernet Addressing

IEEE 802.5 token ring networks (and FDDI) differ from Ethernet networks in the transmission order of bits in each octet. Ethernet nodes always transmit the least-significant bit of a byte first. Token ring nodes always transmit the most-significant bit of a byte first (see figure 6). The problem this poses for bridging is in the interpretation of station addresses. Problems in interpreting station addresses arise in a couple of different ways.

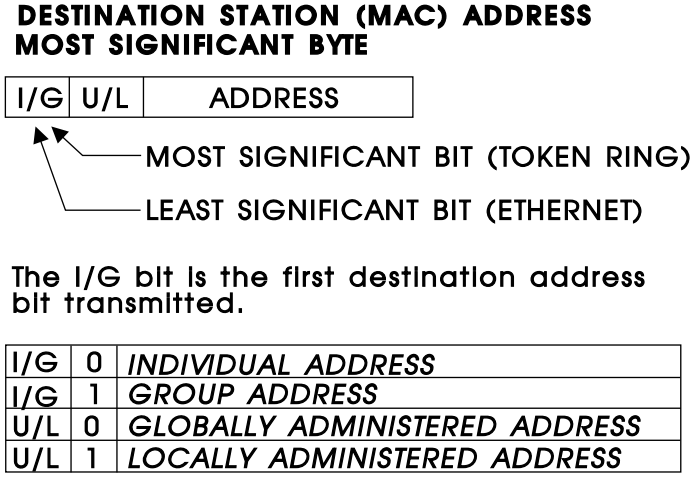

Figure 6. Ethernet/Token Ring Bit Order Differences

First, consider the network in figure 7. Assume station addresses are statically assigned. Assume system A's station address is 080009000000H, and system B's station address is 010000000000. Each address is an individual address on its associated LAN. When packets are transmitted, the first part of the MAC header transmitted is the *destination* station address. When A sends a frame to B, the first bit transmitted is a "1", since the order of transmission for Ethernet is LSB (least-significant bit) of the most-significant byte first. This bit is also the multicast bit. Thus, the packet that was intended to be sent as a unicast packet is unintentionally sent as a multicast packet. Other systems on the Ethernet will interpret the packet transmitted by system A as a multicast packet. The accidental multicast packet may be misinterpreted by systems as a valid multicast packet, with unpredictable results. Worse, all bridges connected to the Ethernet will forward the packet. In this case, normal unicast packets are treated the same as broadcast packets, with potentially disastrous effects on network utilization.

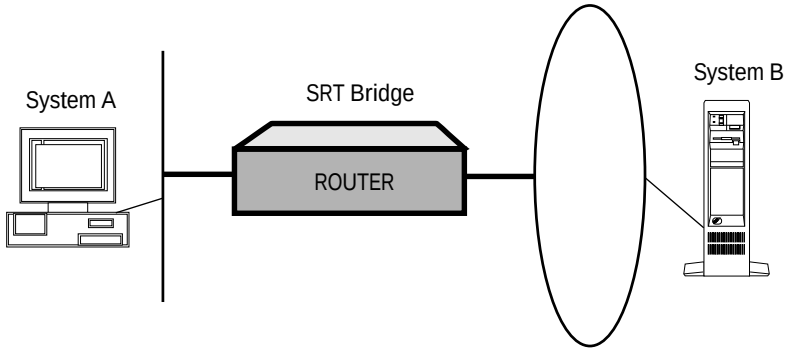

Figure 7. Unsupported Ethernet/Token Ring Bridged Network

The second problem arises when station addresses are determined dynamically. ARP (the Address Resolution Protocol) is used by IP end nodes to determine the station address of a destination system for which the IP address (network-layer address) is known.

Figure 8 shows the sequence of packets sent when system A (from figure 7) tries to determine the station address of system B using ARP, and then sends a data packet to B. First, system A broadcasts an ARP request, which the SRT bridge forwards onto the token ring. The packet is received by system B, which responds by sending an ARP reply with its station address appropriately inserted in the ARP reply packet in the proper order for its media type. Next, system A uses the station address obtained in the ARP reply to send a data packet to system B. This address, however, is improperly ordered for the Ethernet network on which it will now be used, since it was returned (in the ARP reply) by a system on a token ring. This is again the first problem.

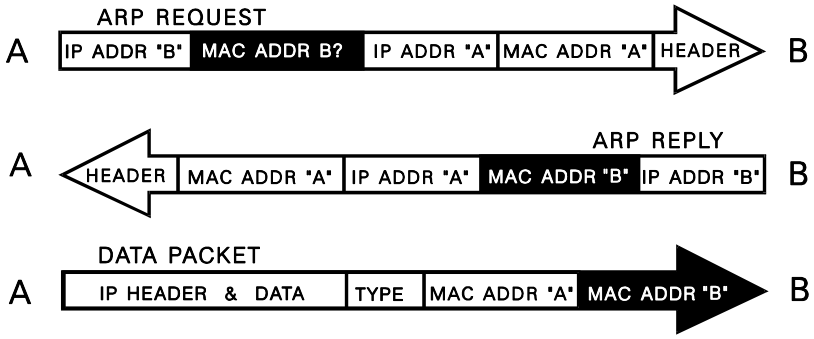

Figure 8. A dynamically determined address is misordered.

Tunneling

Although the SRT bridge does not support bridging data between Ethernet systems and token ring systems for the reasons described above, tunneling provides a mechanism to transmit Ethernet data over token ring backbones and to transmit token ring data over Ethernet backbones. Since source-routing support has not been extended to DECnet or XNS, this capability is especially useful for connecting Ethernet-resident DEC VAXes, XNS systems, and systems using non-routable protocols over token ring backbone networks. Figure 9 shows HP routers connecting Ethernet-resident DEC VAXes and terminal servers over a token ring backbone by using tunneling.

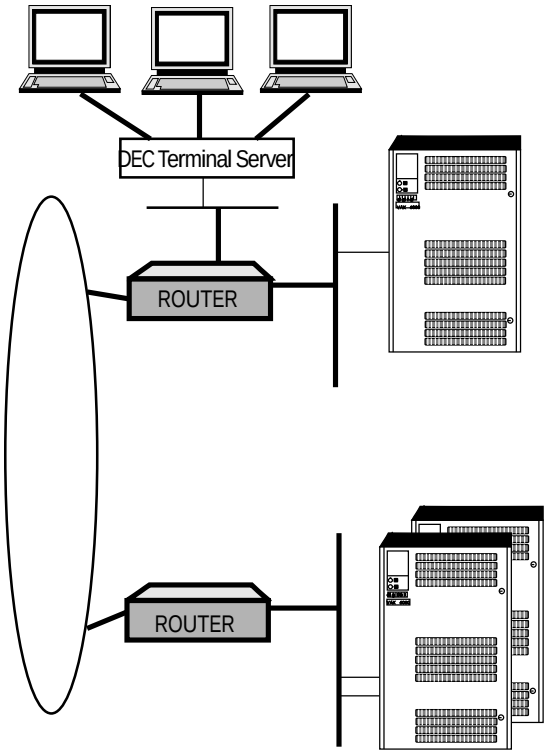

Figure 9. Tunneling Used to Connect DEC VAXes Through a Token Ring Backbone

Tunneling is a feature of the bridgingservice. It is not specifically enabled or configured. Instead, when the bridging software recognizes that an Ethernet packet is being transmitted onto a token ring network, the bridged packet is encapsulated for transmission to another HP router. Thus, tunneling requires one HP router to encapsulate the packet and another HP router to de-encapsulate the packet. Figure 10 shows the encapsulation performed on Ethernet packets to be transmitted over token ring backbones and on token ring packets to be transmitted over Ethernet backbones.

Tunneling Direction:

Ethernet ——— **Token Ring** ——— **Ethernet**

SRT bridge tunnel packet:

Tunnel Header	Ethernet Header	Data

Tunnel Header:

Destination Station Address (48 bits):	Destination SRT Bridge *
Source Station Address (48 bits):	Source SRT Bridge
DSAP (8 bits):	AA (SNAP)
SSAP (8 bits):	AA (SNAP)
Control (8 bits):	03
MSID (24 bits):	000000
Type (16 bits):	8103

Tunneling Direction:

Token Ring ——— **Ethernet** ——— **Token Ring**

SRT bridge tunnel packet:

Tunnel Header	Token Ring Header	Data

Tunnel Header:

Destination Station Address (48 bits):	Destination SRT Bridge *
Source Station Address (48 bits):	Source SRT Bridge
DSAP (8 bits):	AA (SNAP)
SSAP (8 bits):	AA (SNAP)
Control (8 bits):	03
MSID (24 bits):	000000
Type (16 bits):	8101

* ***Note***: If destination SRT bridge address in unknown, address 0100A2FFFFFF is used.

Figure 10. Encapsulation Used to Tunnel Bridge Packets

Limitation Note that when tunneling Ethernet data over a token ring backbone, packets may only traverse HP routers functioning as SRT bridges. Only HP routers and Wellfleet routers understand the tunneling encapsulation.

Source-Routing Bridging

The term source routing was coined by IBM to describe a method of bridging frames across token ring networks. Source-routing bridges differ from transparent bridges in two critical ways:

- Source-routing bridges tolerate a multiplicity of paths between any two nodes in the extended network; transparent bridges, in contrast, require a loop-free topology.
- Source-routing bridges require hosts to supply the information needed to deliver a frame to its intended recipient. Within a source-routing extended network, bridges need not maintain address tables. Rather they make the decision to forward or to drop a frame solely on the basis of data contained within the frame itself. To implement such a scheme, each source node determines the route to a destination node through a process called route discovery.

The route discovery process is enabled by four types of routing directives, listed below. Each type is known by several names, some of which are listed here.

All routes explorer (ARE),
All paths explorer (APE),
All routes broadcast (ARB),
All paths broadcast:

Generates multiple frames that traverse all paths between source and destination stations. Such frames are called all routes explorer (ARE) or all paths explorer (APE) frames, or one of the other names listed above. Upon receiving an ARE frame, each bridge within the extended network appends a routing designator. A routing designator is an information triplet which takes the following form:

[LAN ID i] [Bridge ID] [LAN ID j]

where:

LAN ID i	is a unique number that identifies the LAN (or ring) upon which the ARE frame arrived.
Bridge ID	is a number that identifies the intervening bridge.
LAN ID j	is a unique number that identifies the LAN (or ring) upon which the ARE frame is relayed by the bridge.

After adding a routing designator, each bridge forwards the frame onto all ports except the port on which the frame was received. As a consequence, multiple copies of the same ARE frame can appear on a LAN, and the frame recipient can receive multiple copies of the frame (one copy for each possible path through the extended network). Each ARE frame received by the recipient contains a unique sequenced list of routing designators tracing the frame's path through the extended network.

Spanning tree explorer (STE),
Spanning tree broadcast (STB),
Single route explorer (SRE),
Single path explorer (SPE),
Single route broadcast (SRB),
Single path broadcast:

Generates a single frame that follows a loop-free (spanning-tree-derived) path from source node to destination node. Such frames are called single route explorer (SRE) or transparent spanning frames (TSF), or one of the other names listed above. Upon receiving an SRE, each bridge on the spanning tree forwards the frame onto all active (non-blocked) ports except the port on which the frame was received. With spanning tree broadcast routing, one copy of the SRE appears on each LAN, and the frame recipient receives only a single copy of the frame.

Some other terms referring to a looped topology are alternate routes, parallel bridges, and redundant bridges.

Specific routing:

Generates a single frame that traverses a specific path designated by the source node. Such a frame is called a specifically routed frame (SRF). SRFs contain a list of routing designators that maps a unique path through the extended network from source to destination node. Upon receiving an SRF, each bridge examines the list of routing designators. It forwards the SRF only if it is on the specified path, otherwise it ignores the frame.

Null routing:

Indicates that the source node does not desire any routing services from network bridges. As a result, null-routed frames are restricted to the resident LAN of the originating node.

How Source Routing Works

Source routing networks consist of LAN segments interconnected by source routing bridges. Each LAN segment has an identification number unique throughout the network, called a LAN ID, and also called a ring number or ring ID. Each source routing bridge has an identification number; the source routing bridge is always (by default) bridge ID number 1. Additionally, each source-routing bridge has an internal LAN ID number unique throughout the network.

As a source-routed frame traverses the network, it collects a sequence of routing designators that track its path through the network. Every source-routing bridge that the frame passes through inserts routing designators in the frame's MAC header. Transparent bridges, however, do not write to the MAC header. Each routing designator pairs a LAN segment number with a bridge ID number in order to identify a portion of the frame's path through the bridge.

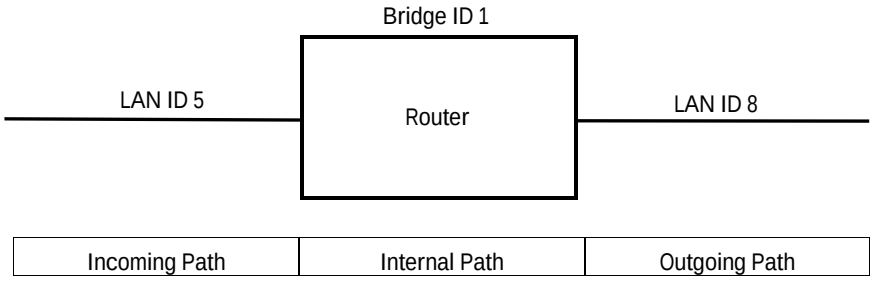

Figure 11. SRT Bridge Routing Designators

In figure 11, the router receives the source-routed frame on LAN segment 5 and relays the frame on LAN segment 8. Consequently, the router adds three routing designators (5-1, A-1, 8-0) to the frame's MAC header. Figure 12 illustrates how these routing designators map the frame's path through this router.

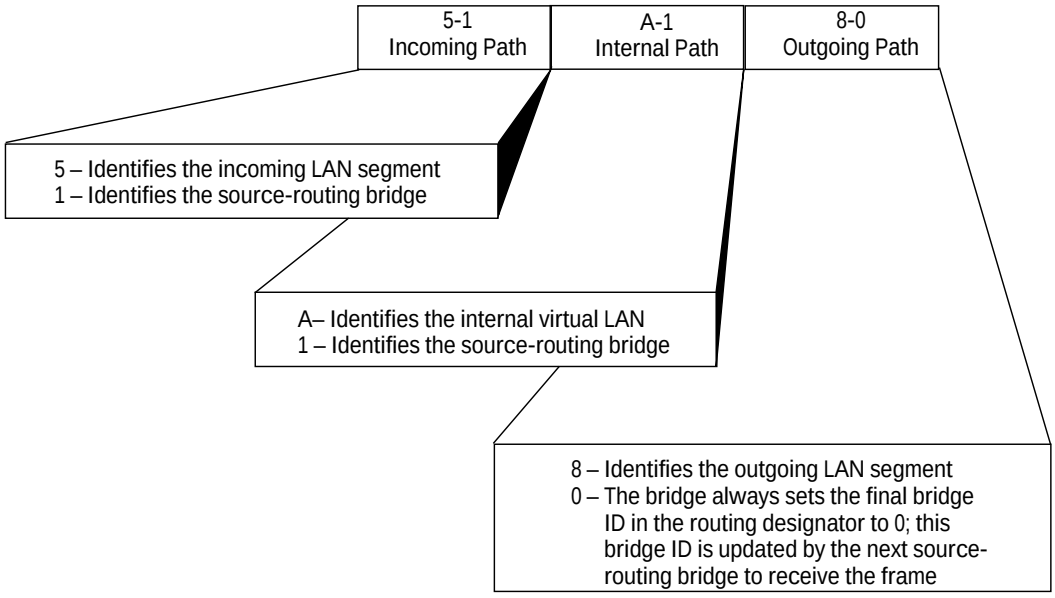

Figure 12. Sample Routing Designators

When a source-routed frame reaches its destination, its MAC header contains the route (identified by a sequence of routing designators) that the frame traversed to get there.

For example, figure 13 depicts a multi-ring network with three SRT bridging routers (A, B, and C), and one transparent bridge (T). Node H1 wants to use source routing to exchange frames with node H2. In order to initiate route discovery, H1 transmits an ARE frame addressed to H2. As the ARE frame crosses the network, every source-routing bridge that it passes through inserts routing designators in the frame's MAC header. If the frame crosses a transparent bridge, the bridge simply forwards the frame on the basis of its destination station address.

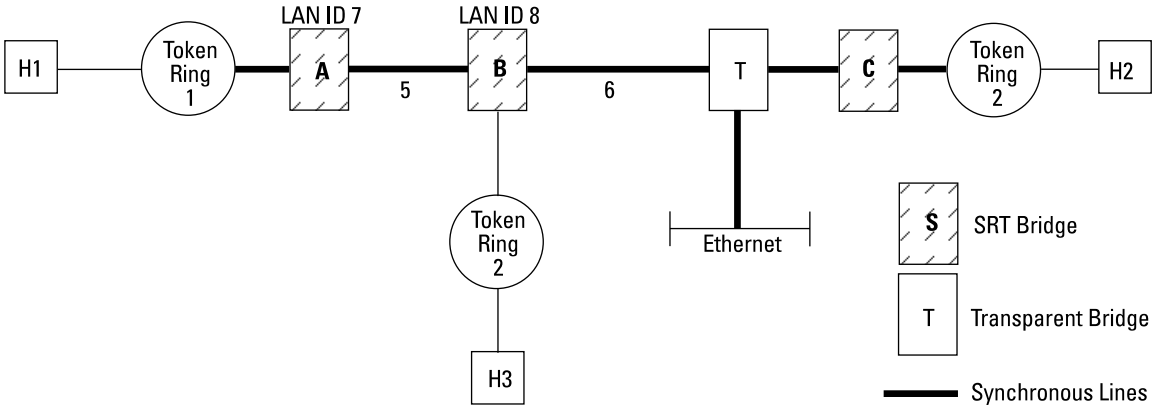

Figure 13. Multi-Source-Routed Network

In figure 13, router A is the first source-routing bridge to receive the ARE frame originated by H1. A inserts its routing designators (1–1, 7–1, 5–0) in the frame's MAC header:

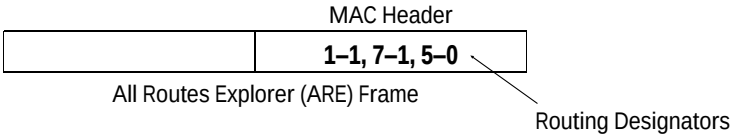

Figures 14, 15, and 16 depict how the routing-designators (1–1, 7–1, 5–0) identify the frame's path through bridge A. Note that in figure 14, bridge A inserts the incoming LAN ID/bridge ID pairing (1–1) only because it is the first source-routing bridge to receive the frame. If A was not the first source-routing bridge to receive this frame, the incoming LAN ID/bridge ID pairing would be taken from the outgoing LAN ID/bridge ID pairing inserted by the previous source-routing bridge.

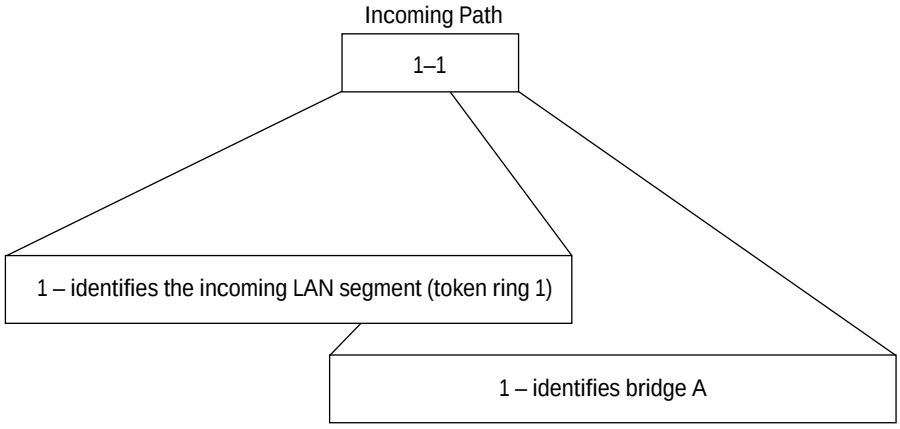

Figure 14. Routing Designator 1

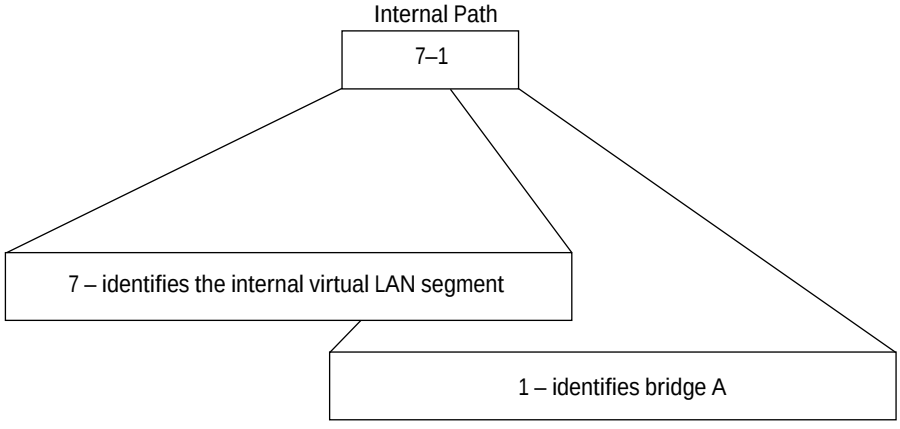

Figure 15. Routing Designator 2

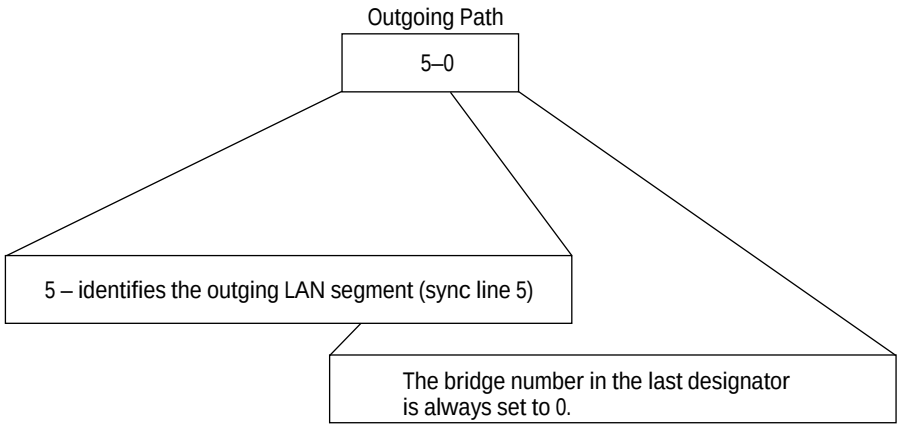

Figure 16. Routing Designator 3

Router B (in figure 13) is the next source-routing bridge to receive the frame. B updates A's last routing designator by changing bridge ID 0 to 1, and then inserts the remainder of its routing designators (8–1, 6–0) in the frame's MAC header. B then floods ARE frames onto LAN segments 3 and 6 with the following routing designators:

	1–1, 7–1, 5–1, 8–1, 3–0

The above ARE frame is transmitted onto LAN segment 3, where the frame is eventually dropped, since LAN segment 3 does not provide a path to H2.

	1–1, 7–1, 5–1, 8–1, 6–0

The above ARE frame is transmitted onto LAN segment 6; this frame passes through transparent bridge T.

Because T (in figure 13) is a transparent bridge, it does not recognize the routing designators in the source-routed frame. T treats the ARE frame as if it was a transparent-bridging frame, and simply forwards the frame to router C and the Ethernet (the frame transmitted onto the Ethernet is discarded).

Router C (in figure 13) is the last source-routing bridge to receive the ARE frame. C updates B's last routing designator by changing bridge ID 0 to 1, and then inserting the remainder of its routing designators (9–1, 2–0) in the frame's MAC header. C then transmits the ARE frame to the destination node H2.

H2 inspects the frame's MAC header to learn the route that this particular frame traversed. The route is mapped out by the sequence of routing designators (1–1, 7–1, 5–1, 8–1, 6–1, 9–1, 2–0) inserted by source-routing bridges A, B, and C.

Note

Other route discovery protocols are available and extensively used. All such protocols, however, exchange TSFs, ARE frames, and SRFs.

Source-Routing Bridging on HP Routers

This section describes the routing of frames through a network using HP's source-routing architecture. The HP router configured as a source-routing bridge handles incoming packets differently depending on its position in the network. The sample network shown below in figure 17 contains bridging HP routers and two end nodes (H1 and H2). The internal LAN ID, group LAN ID, and bridge ID are listed for each bridge in column below the bridge illustrated. Ring 1 has a LAN ID of 1; ring 2 has a LAN ID of 2; ring 3 has a LAN ID of 3; ring 4 has a LAN ID of 4; ring 5 has a LAN ID of 5. This is the sample network used in the subsequent illustrations throughout this section, figures 18 through 23. An HP router that is performing source-routing bridging is referred to as a "bridge" in this section.

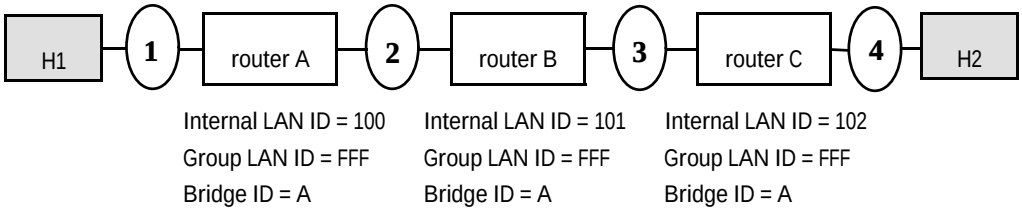

Figure 17. Sample HP Source-Routing Bridging Network

The following three subsections describe three tracks, listed below, taken by source-routing frames in the sample network shown in figure 17. The subsection for each track gives different scenarios for bridges in the different network positions—first, middle, and last—describing how they manipulate the frame's Routing Information Field (RIF).

1. Track the RIF of an explorer frame (ARE) sent from H1 to H2.

2. Track the RIF of the specifically routed frame (SRF) sent back from H2 to H1.

3. Track the RIF of a specifically routed frame (SRF) sent from H1 to H2.

Track 1. An Explorer Frame From Node 1 to Node 2

This section tracks explorer frames (AREs) sent from H1 to H2 in the sample HP source-routing bridging network. Figure 18 below illustrates the same network as does figure 17, except that arrows in figure 18 indicate the direction of the frame's path.

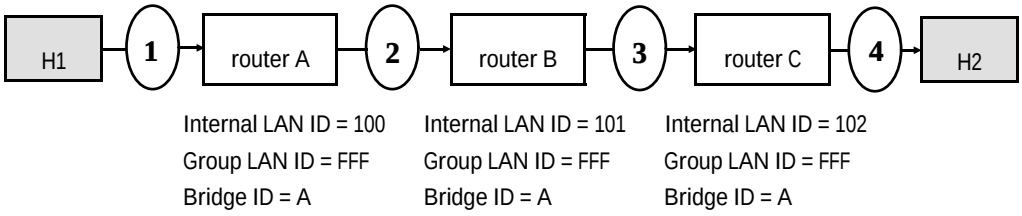

Figure 18. Tracking an Explorer Frame

First bridge receiving the explorer frame: The explorer frame received by bridge A from ring 1 has not traversed any other bridges. A simply does the following to the RIF before transmitting the frame toward ring 2. (Figure 19 illustrates the RIF in the frame.)

- Adds the incoming LAN ID, and its bridge ID
- Adds the internal LAN ID, and its bridge ID
- Adds the outgoing LAN ID, and bridge ID of 0

Other bridges receiving the explorer frame: The explorer frame received by bridges B and C contains internal LAN IDs (thus indicating that this frame has traversed at least one other bridge). Each of these bridges does the following to the RIF before transmitting the frame toward rings 3 and 4. (Figure 19 illustrates the RIF in the frame.)

- Keeps the outgoing LAN ID as its incoming LAN ID; replaces the bridge ID of 0 with its own bridge ID.
- Replaces the internal LAN ID and bridge ID of the last bridge traversed with its own internal LAN ID and bridge ID.
- Adds the outgoing LAN ID, and bridge ID of 0.

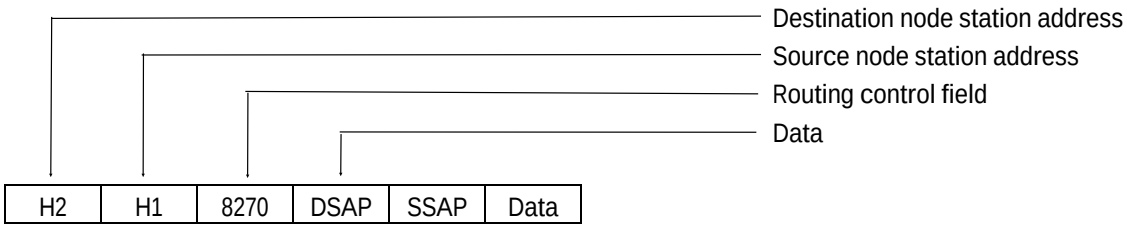

Frame received by bridge A

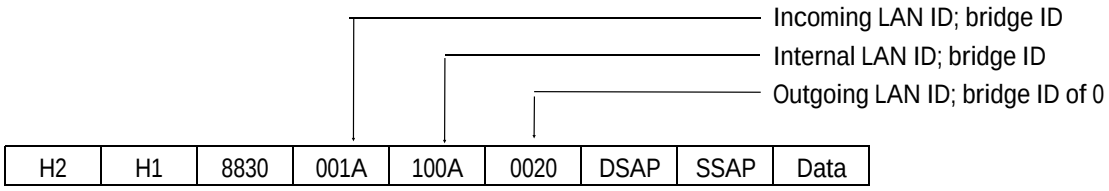

Frame sent by bridge A onto ring 2

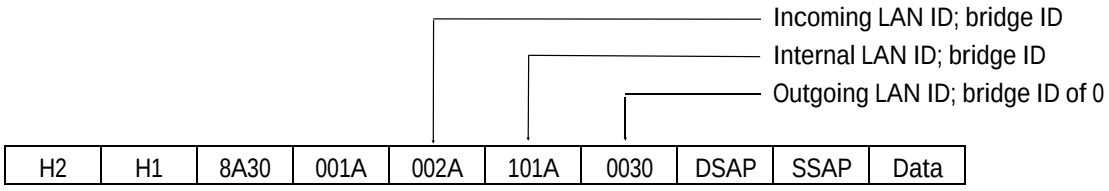

Frame sent by bridge B onto ring 3

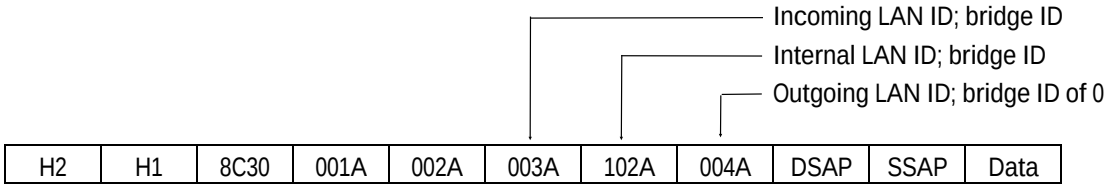

Frame sent by bridge C onto ring 4

Figure 19. Structure of an Explorer Frame

Track 2. The Specifically Routed Frame Back to Node 1

This section tracks specifically routed frames (SRFs) sent back from H2 to H1. See figure 20.

If there is only a single bridge ID (for an HP router) in the RIF, then the bridge simply transmits the frame to the outgoing circuit group without making any modification. This is only true when the frame only has to traverse a single bridge between the source and destination end nodes. Because of the simplicity of this case, it is not described in any further detail here.

On the other hand, if there are multiple bridge IDs (for an HP router) in the RIF, the RIF is manipulated differently by the bridges in the first, middle, and last positions—which are C, B, and A, in that order, in the direction taken by this frame.

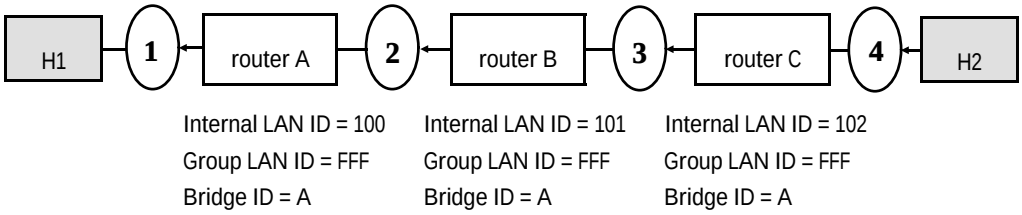

Figure 20. Tracking a Specifically Routed Frame Back to Node 1

First of several bridges receiving the SRF: The frame received by bridge C from ring 4 has not traversed any other bridges yet. (But there are multiple bridge IDs in the RIF, so it will.) This bridge does the following to the RIF before transmitting the frame toward ring 3. (Figure 21 illustrates the RIF in the frame.)

- Changes the destination node's station address at the beginning of the frame to an HP group address. This address appears as C000A2FFFFF*x*, where *x* is the bridge ID of the next bridge specified by the RIF.
- Removes its own internal LAN ID.
- Inserts the group LAN ID before the last incoming LAN ID/bridge ID listed in the RIF. (See figure 21.) Eventually, the group LAN ID will be replaced with the internal LAN ID of the last bridge along the frame's path.
- Copies the destination node's station address into the data portion of the frame.

Between the first and last bridge receiving the SRF: The frame received by bridge B from ring 3 has traversed at least one other bridge. However, this is not the last bridge that the frame must traverse. This bridge does the following to the RIF before transmitting the frame toward ring 2. (See figure 21.)

- Locates the bridge ID that is located at the end of the HP group address.
- Changes the bridge ID at the end of the HP group address to the bridge ID of the next bridge in the RIF. (In this example, all bridge IDs are the same, so the frame is not modified.)

Last of several bridges receiving the SRF: Bridge A is the last of several bridges traversed by the frame. This bridge does the following to the RIF of the frame it receives, before transmitting the frame toward ring 1. (See figure 21.)

- Replaces the HP group address with the destination station address that was saved to the data field.
- Replaces the group LAN ID with its own internal LAN ID.

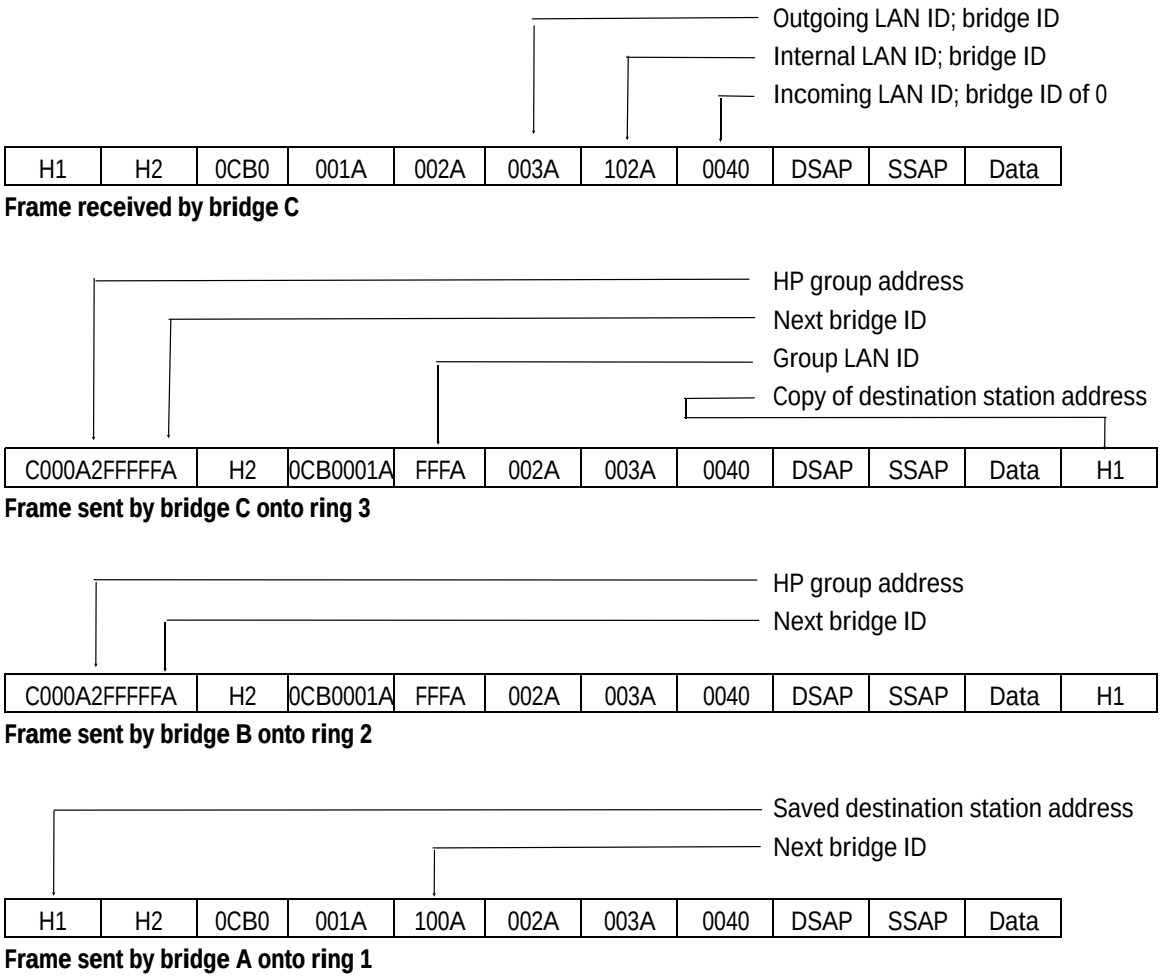

Figure 21. Structure of a Specifically Routed Frame Back to Node 1

Track 3. A Specifically Routed Frame From Node 1 to Node 2

This section tracks specifically routed frames (SRFs) from H1 to H2. HP's source routing algorithm works the same as when H2 routes a specifically routed frame to H1. To follow this case, refer to the sample network illustration in figure 22.

If there is only a single bridge ID (for an HP router) in the RIF, then the bridge simply transmits the frame to the outgoing circuit group without making any modification. This is only true when the frame only has to traverse a single bridge between the source and destination end nodes. This case is not described in any further detail here.

On the other hand, if there are multiple bridge IDs (for an HP router) in the RIF, the RIF is manipulated differently by the bridges in the first, middle, and last positions—A, B, and C, respectively, for this case.

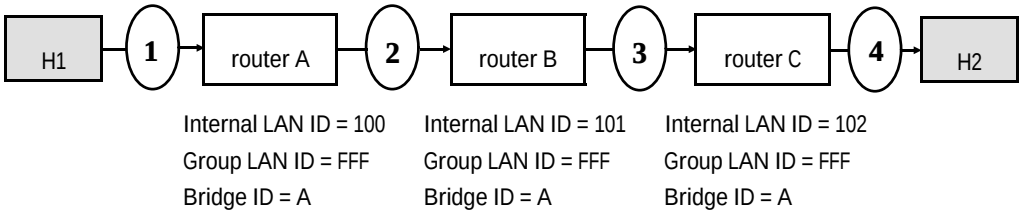

Figure 22. Tracking a Specifically Routed Frame from Node 1 to Node 2

First of several bridges receiving the SRF: Bridge A is the first to receive the frame from ring 1 and does the following to the RIF before transmitting the frame toward ring 2. (Figure 23 illustrates the RIF in the frame.)

- Changes the destination node's station address at the beginning of the frame to an HP group address. This address appears as C000A2FFFFF*x*, where *x* is the bridge ID of the next bridge specified by the RIF.
- Removes its own internal LAN ID.
- Inserts the group LAN ID before the last Incoming LAN ID /bridge ID listed in the RIF. (See figure 23.) Eventually, the group LAN ID will be replaced with the internal LAN ID of the last bridge along the frame's path.
- Copies the destination node's station address into the data portion of the frame.

Between the first and last bridge receiving the SRF: The next bridge B does the following to the RIF before transmitting the frame toward ring 3. (See figure 23.)

- Locates the bridge ID that is located at the end of the HP group address.
- Changes the bridge ID at the end of the HP group address to the bridge ID of the next bridge in the RIF. (In this example, all bridge IDs are the same, so the frame is not modified.)

Last of several bridges receiving the SRF: The last bridge to receive the frame does the following to the RIF before transmitting the frame toward ring 1. (See figure 23.)

- Replaces the HP group address with the destination station address that was saved to the data field.
- Replaces the group LAN ID with its own internal LAN ID.

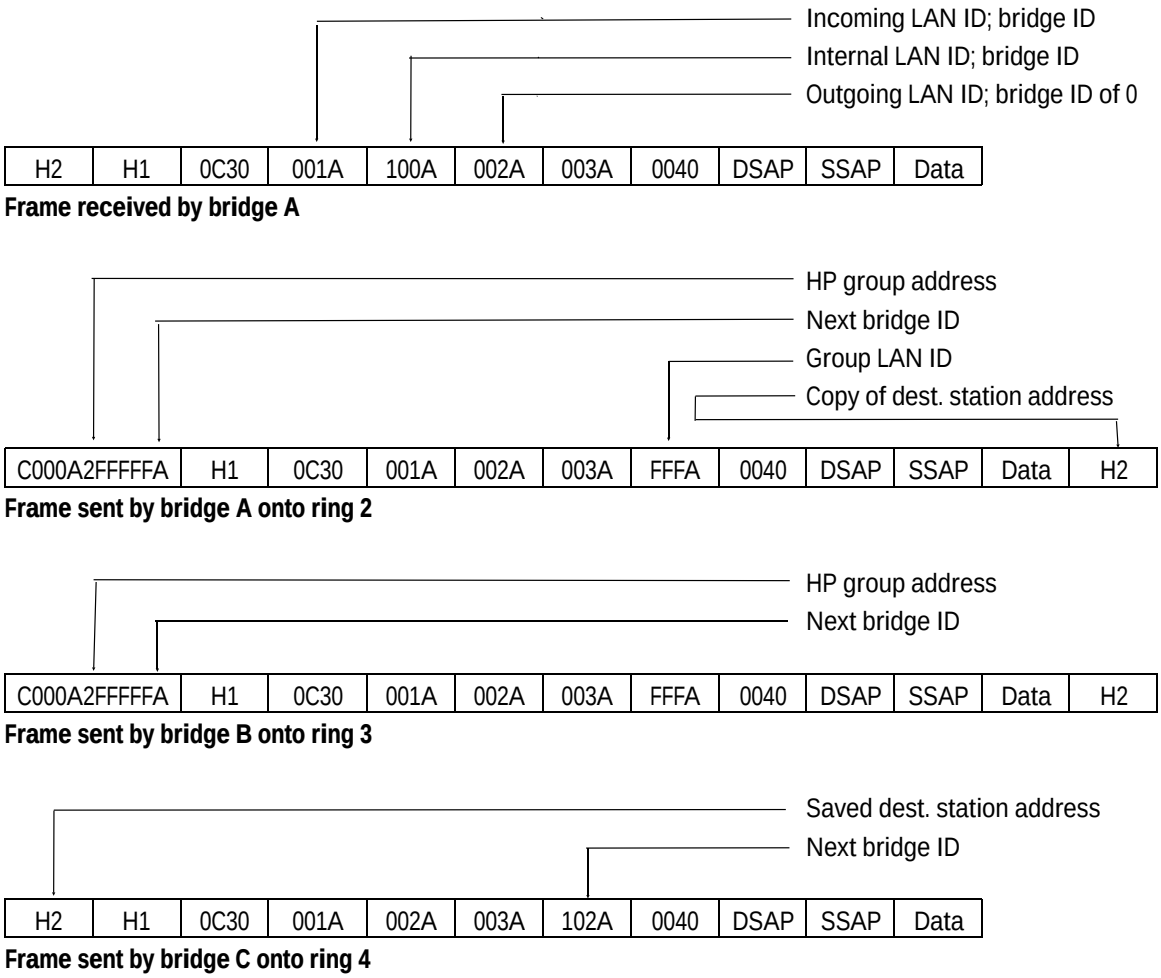

Figure 23. Structure of a Specifically Routed Frame from Node 1 to Node 2

Source-Routing/Transparent Bridging

The SRT bridging provides concurrent transparent and source-routing services. Figure 24 shows a sample multi-ring, multi-Ethernet extended network linked by four routers serving as SRT bridges. Router T provides only transparent bridging services. The three other routers (all labeled S) provide both source-routing and transparent bridging services, when they have source routing enabled. However, HP routers do not support source-routing transparent bridging to traverse both Ethernet and token ring LANs using a single bridging technique. The two different methodologies are provided concurrently.

The transparent bridge treats all frames as if they are transparent-bridging frames. In order to effect route discovery, however, the SRT bridge needs to separate frames which require source-routing service from those frames which require transparent-bridging service.

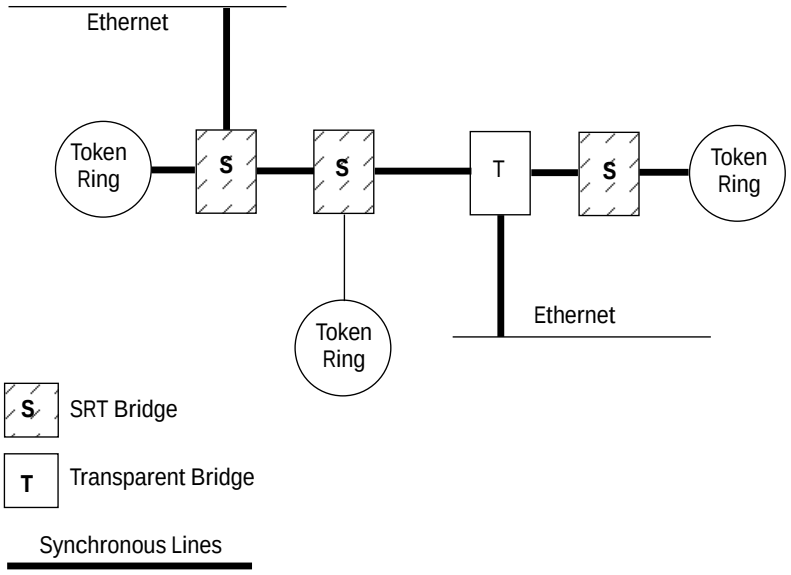

Figure 24. Sample SRT Topology

In order to identify source-routing frames, the SRT bridge inspects the value of the most significant bit of the frame's source address (referred to as the routing information indicator or RII). An RII value of 1 specifies source routing; an RII value of 0 specifies transparent bridging.

Source Route Translational Bridging (TRNSB)

Source Route Translational Bridging (TRNSB) translates frames between source-routing bridging (SRB) circuit groups and transparent bridging (TB) circuit groups. The router translates frames for protocols such as SNA or NetBEUI between token ring circuit groups configured for SRB and Ethernet circuit groups configured for TB. (NetBEUI is the name of the protocol driver installed as the default protocol for MS LanManager and IBM LanServer). TRNSB also converts for WAN links configured for SRB or TB, as appropriate.

Note

Before reading on, ensure that you understand the explanation of transparent bridging (TB), source-routing bridging (SRB), and global source routing, earlier in this section.

Figures 25 through 28 below illustrate some common topologies that can use TRNSB.

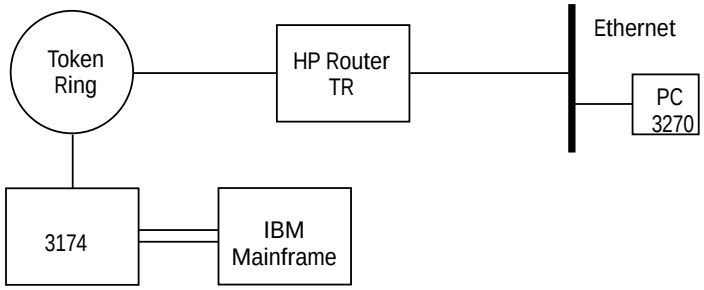

Figure 25. Simple One-Hop TRNSB Topology

In figure 25, the PC running the PC3270 application is limited to either 802.3 or Ethernet frame format, and transparent bridging. (SRB is not supported on 802.3 or Ethernet.) The 3174 is limited to 802.5 framing, and only supports source route bridging. The HP Router TR will provide the necessary conversion for these devices to communicate.

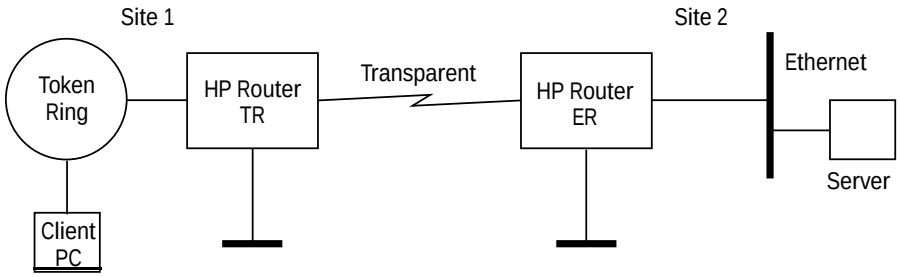

Figure 26. Transparent WAN Backbone Topology

Figure 26 illustrates a common WAN topology that can use TRNSB. In this case, the primary bridging technology is transparent bridging. Thus, the WAN circuit group is not configured for source-routing bridging. Most traffic crossing the WAN is TB traffic going to and from the Ethernets. The user also has a need to print to the HP LaserJet printer on the Ethernet at site 2 from the client PC at site 1. The HP Router TR at site 1 will convert the NetBEUI/SRB/802.5 frames to NetBEUI/TB/802.3 frames.

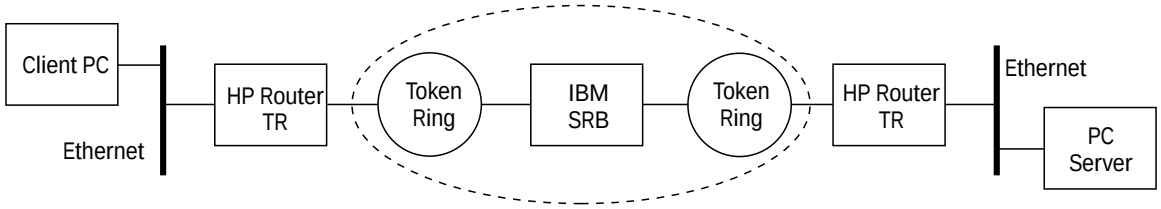

Figure 27. Token Ring Backbone Topology

Figure 27 shows an example of a 16 Mbit/s token ring SRB backbone, used to interconnect two distributed Ethernets. This topology is often found in environments where many departments have installed their own LANs, while the site services organization (often in coordination with the information technology organization) has specified token ring as the backbone technology.

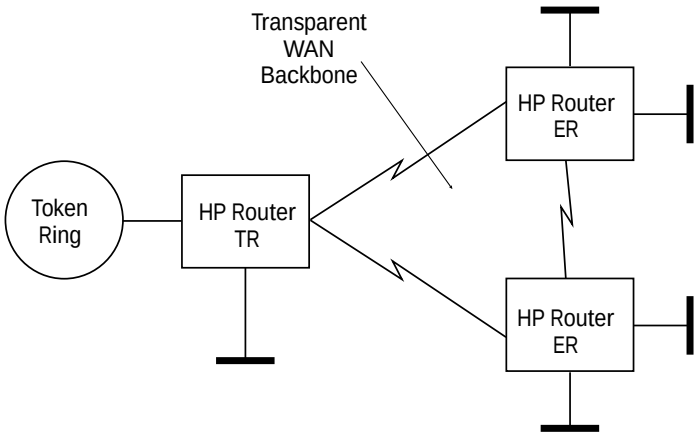

Figure 28. Transparent Meshed WAN Backbone Topology

Finally, figure 28 shows how TRNSB could be used in a meshed and transparently bridged WAN backbone. To the TRNSB, this topology is really not different from that of figure 26. However, it demonstrates how TRNSB can be used in a spanning tree environment. If the network in figure 28 had mostly token ring SRB networks, the WAN circuit groups would more likely be configured as source-routing links. This would work also. Please note that with TRNSB functionality enabled, the WAN links must all be configured as source routing, or all not source routing, but cannot be mixed.

There are two basic functions of the TRNSB, frame format conversion and bridge technology conversion, as detailed below.

Frame Conversion

Frame format conversion is merely moving the fields of the Medium Access Control (MAC) layer headers, or creating or deleting them. Figures 29 and 30 show the frame format conversions performed by the TRNSB.

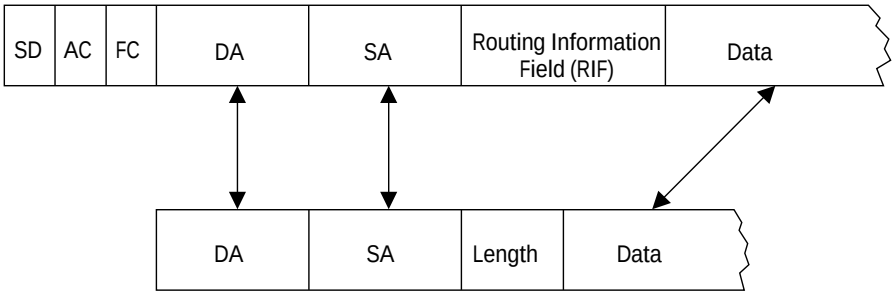

Figure 29. 802.5 to 802.3 Frame Conversion

Figure 29 shows how 802.5 frames are converted to and from 802.3 frames. Since bit ordering differs between these types of networks, the bits of the station addresses (DA, SA) must be reversed. The Routing Information Field (RIF) is extracted when moving from 802.5 to 802.3, and inserted when going in the other direction. (The TRNSB maintains a store of RIFs; see below.) The 802.3 length field is also created or removed as necessary.

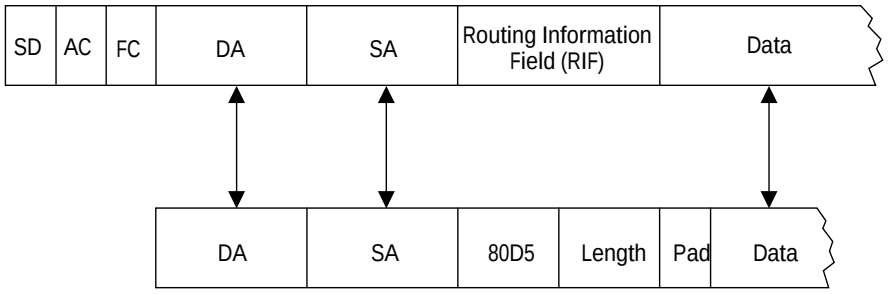

Figure 30. 802.5 to Ethernet Frame Conversion

Figure 30 shows conversion to the Ethernet version 2 frame format. (This is often called the PC/RT format, or 80D5 format.) The conversion is similar to that of figure 29.

Bridge Technology Conversion

From the perspective of the nodes on the Ethernet, the bridge technology conversion is an algorithm that makes it appear that all nodes on the token-ring SRB side are running on the same Ethernet LAN. From the perspective of the nodes on the token ring, this algorithm also performs the opposite function, of making the Ethernet nodes appear to be performing source routing.

To understand how the TRNSB converts between the bridging technologies, begin by examining the conversion from SRB to TB. There are two basic types of SRB frames:

- Explorer frames
- Specifically routed frames

When the router receives an explorer frame from an SRB circuit group, the TRNSB code accesses the address table used for transparent bridging to see if the node is known to be on one of the transparent bridging circuit groups. If it is found, the frame is converted to Ethernet/802.3 format and forwarded to that circuit group. If the destination is not known, then the TRNSB converts the frame and floods it to all transparent circuit groups configured for TRNSB. If the destination address in the frame indicates a functional address or the broadcast address, the frame is always flooded. In all of these cases, the source of the frame is learned in the address table. In addition to learning the association of a station address with a circuit group, the Routing Information Field (RIF) for that station is stored. (It will be used when converting in the other direction.)

When the router receives a specifically routed frame, it converts the frame only if the next hop in the Routing Information Field (RIF) indicates that a transparent circuit group is the next hop in the source routing path.

When converting from TB circuit groups to SRB circuit groups, the TRNSB converts to a specifically routed frame if the destination station address is found in the address table. The RIF that was stored (as described above) is inserted into the frame, and the frame is routed through a source-routed network.

When the destination is not found in the address table or when the destination is a multicast address, the frame is converted to an explorer frame (single route explorer). This frame is then forwarded to all SRB circuit groups configured for TRNSB.

As is shown in figures 29 and 30, both 802.3 and Ethernet formats are supported. Ethernet version 2 frames are limited to what is often called the IBM PC/RT format, which is a proprietary encapsulation of 802.2 frames in an Ethernet type header. Which of the 802.3 and Ethernet formats is used is indicated by configuration of the router. Either one can be configured as the default conversion format, and a list of station addresses that use the other format can be specified. Most end nodes use the 802.3 format.

Limitations

- TRNSB does not support protocols that can be routed.
- TRNSB does not operate with looped topologies that mix SRB and TB. That is, TRNSB operates in looped topologies only when the links in the loops are either all SRB or TB, but not a combination of both.
- TRNSB does not support either the hop-count reduction or the group LAN ID function described in the following sections of this chapter.
- The HP Remote Bridge will not perform the TRNSB translation, but can be connected across a transparent-bridging WAN link to an HP router that does perform the translation—an HP Router TR or HP Router 650 with a token ring interface module.
- The load-balancing function for bridging circuit groups does not support translational bridging.

Configurable Hop-Count Reduction Algorithm

In a source-routing bridging environment, a frame is generally limited to seven bridge hops. This standard source-routing algorithm counts one hop per intervening bridge and imposes a limit of seven hops from source-routing source to source-routing destination. A special HP alternative source-routing-bridging algorithm called hop count reduction bypasses this limit by providing an infinite hop capability. This algorithm offers the advantage of supporting larger diameter networks than were previously possible.

The Hop Count Reduction parameter determines which algorithm is used. (See the Hop Count Reduction parameter description in the configuration reference documentation supplied with your router.) The standard source-routing algorithm is in effect by default. Thus, to allow an infinite number of hops, set Hop Count Reduction to Yes. But if your network never needs more than seven hops, it is recommended that you leave this parameter set to the default No.

Note that, when used with certain network topologies, the hop count reduction algorithm generates extra explorer frame traffic and multiple redundant route responses. In certain topologies, such redundant traffic may be insignificant. In others, end nodes and/or the bridging router may experience difficulty in detecting loops or routes.

Note

Unless the network diameter is greater than seven hops, it is recommended that the hop count reduction algorithm be configured off, so that the standard seven-hop source-routing algorithm is used. This minimizes router selection paths and redundant explorer traffic. In topologies requiring hop-count reduction, you should make efforts to reduce physical loops around the bridge.

Do not configure the hop count reduction algorithm with source route translational bridging (TRNSB).

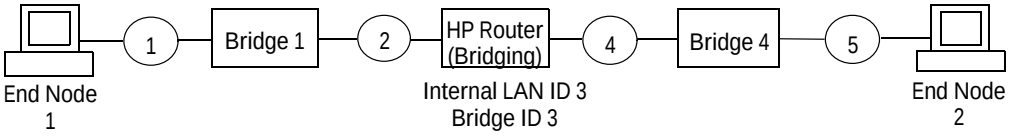

Figure 31. Sample Network for Hop Count

Hop Count Reduction Algorithm Enabled

In figure 31 (above), with Hop Count Reduction set to Yes, end node 1 sends an explorer frame. The RIF at ring 1 is 8270; the RIF at ring 2 is 8670 0011 0020.

The bridging HP router, with internal LAN ID 3 and bridge ID 3, receives the frame from ring 2 and saves this RIF as the path to end node 1. It then removes the RIF from the frame and inserts its internal LAN ID and the outgoing LAN ID. The RIF at ring 4 is then 8670 0033 0040; the RIF at ring 5 is 8870 0033 0044 0050.

End node 2 received the explorer frame and saved the RIF as the path to end node 1. End node 2 then responds with a specifically routed frame. The RIF at ring 4 is then 88F0 0033 0044 0050.

The bridging HP router receives this frame on ring 4. The RIF is saved as the path to end node 2. The RIF is removed from the frame and replaced with the RIF saved as the path to end node 1. The RIF on ring 2 of the frame heading back to end node 1 is then 88F0 0011 0023 0030. End node 1 saves this RIF as the path to end node 2.

Hop Count Reduction Algorithm Disabled

In figure 31 (above), with Hop Count Reduction set to the default No, end node 1 sends an explorer frame. The RIF at ring 1 is 8270; the RIF at ring 2 is 8670 0011 0020.

The bridging HP router receives the frame from ring 2 and learns that to get from the router to ring 1 through bridge 1, it needs to send the frame out through ring 2. The router then overwrites the routing designator of the incoming ring, in this case ring 2, with its internal LAN ID, adds the outgoing ring, and then sends the frame out to ring 4. The RIF at ring 4 is 8870 0011 0033 0040, and the RIF at ring 5 is 8A70 0011 0033 0044 0050.

End node 2 receives the explorer frame and saves the RIF as the path to end node 1. End node 2 responds with a specifcally routed frame. The RIF at ring 4 is then 8AF0 0011 0033 0044 0050.

The HP router receives this frame from ring 4 and learns that to get from the router to ring 5 through bridge 4, it needs to send the frame through ring 4. The router overwrites the internal LAN ID in the RIF with the outgoing ring that it learned when the frame first came through. It then overwrites the incoming ring, in the case ring 4, with the internal LAN ID and sends it through ring 2. The RIF at ring 2 is then 8AF0 0011 0023 0034 0050. End node 1 saves this RIF as the path to end node 2.

Partial, Full, and No Reduction

The setting of the Bridge ID configuration parameter affects how the two source-routing algorithms record hops. Both the standard algorithm and HP's hop count reduction algorithm work with the Routing Information Field (RIF) in a source-routed frame's MAC header. The RIF is the part of the frame containing the path between source and destination end nodes. (For more information on bridge IDs, see an earlier section of this note, "How Source Routing Works" and the next section on configuration.)

In a network with multiple bridges, when Hop Count Reduction is off and each router's bridge ID is different, there is a routing designator pair for every LAN or ring (incoming and outgoing path) and for every bridge (an internal path identifying a virtual LAN and the bridge). All the bridges use up RIF space. There is no hop count reduction.

With consecutive bridging HP routers, when Hop Count Reduction is off and each router's bridge ID is the same, then only the first of the bridges uses up RIF space with its virtual, internal LAN ID. This is partial hop count reduction.

With consecutive bridging HP routers, with Hop Count Reduction set to Yes, whether the bridge IDs are identical or not makes no difference. None of the internal LANs for bridges take up RIF space. In addition, all routing designators for the actual LANs or rings between the consecutive bridges are discarded also. This is total hop count reduction.

Configuring Source-Routing Bridging

The parameters commonly used for configuring source routing globally for the bridging service include the following:

- Internal LAN ID
- Bridge ID
- Source Route Bridge ID
- Hop Count Reduction
- Group LAN ID
- Translational Bridge

The parameters commonly used to configure source routing for each individual bridging circuit group include the following:

- Src Rte
- LAN ID
- Block STE
- Translational Bridge

Internal LAN ID On multiport bridges (such as HP's bridging router), the internal LAN ID assigns a numeric identifier to the bridge, which is used in constructing routing designators. The internal LAN ID is also known as a virtual LAN ID. Each internal LAN ID must be unique among all internal LAN IDs and LAN IDs throughout the network, and be different from the group LAN ID.

Bridge ID The bridge ID for this router identifies a specific source-routing bridge (this router). To keep the hop count partially reduced (see the preceding section on hop count reduction, it should generally match the bridge ID assigned to all other HP (or Wellfleet) routers on the network. However, the bridge ID or IDs you assign to HP routers must be different than those used by any other bridge on the network.

If Hop Count Reduction is enabled, then the bridge IDs for the bridging routers can be either the same or different.

If translational bridging (TRNSB) is being used on this router, then you must assign it a unique bridge ID.

If two or more HP routers operate as bridges in parallel, then, to avoid looping traffic, you must assign them different, unique bridge IDs. In this case, you must use the "Source Route Bridge ID" parameter to inform this bridging router of each other HP bridging routers (that is, of the other bridge IDs) that exist on the network.

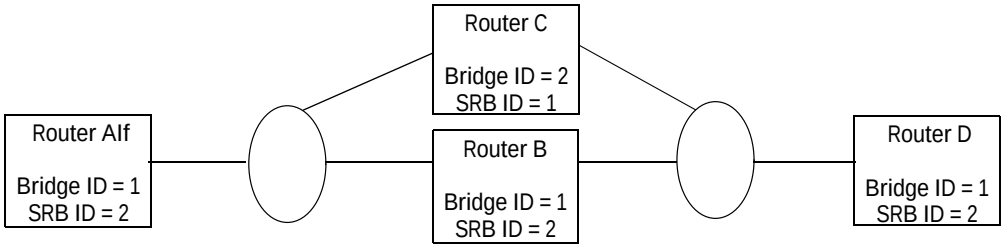

Figure 32. Source-Routing Bridges Operating in Parallel

Figure 32 illustrates an example of parallel (redundant) bridging. Because B and C operate in parallel, the bridge ID for C is changed from 1 to 2, so that it is unique. Because all HP routers have to know about all other HP routers, each bridging service is configured with a "Source Route Bridge ID" (SRB ID) for each other bridge ID on the network. That is, the SRB ID for A, B, and D is set to 2, which specifies that an additional path exists through router C. The SRB ID for C is set to 1, which specifies that an additional path exists through router B.

Hop Count Reduction Configure it to be No unless you want infinite hop capability. No specifies a maximum of seven hops between the source-routing source and destination.

Group LAN ID This is used by the bridging router when transmitting specifically routed frames (SRFs) between HP (or Wellfleet) bridging routers. Together with the other routing designators, the group LAN ID helps intermediate bridges identify the destination end node. The group LAN ID you configure must match the group LAN ID assigned to all HP or Wellfleet bridges and routers, and must differ from all internal LAN ID values and external LAN ID values assigned to any bridges or routers on the network.

Source Route A configuration parameter named "Src Rte" is used to enable source-routing bridging on a circuit group.

LAN ID A number unique throughout the internetwork is assigned to the Ethernet/802.3 LAN or token ring on each circuit group participating in source-routing bridging. The LAN ID is also known as a ring number or ring ID.

Block Spanning Tree Explorer Frames Single-route explorer frames, also called spanning tree explorer (STE) frames or single-route broadcast frames) received on a specific circuit group can be configured to be dropped by changing the “Block STE” parameter for that circuit group to Yes. Doing so does not stop single-route explorer frames from being transmitted on the circuit group.

Translational Bridge This configuration parameter is used to enable TRNSB for each circuit group that will translate frames between source-routing bridging circuit groups and transparent bridging circuit groups.

Another parameter screen for the bridging service on this router (covering all circuit groups) has some parameters specifically for translational bridging. The aging timer can be left at the default. The default frame conversion type can be configured to either 802.3 or Ethernet. Then this default can be varied for each node that requires the other frame conversion type; the nodes are specified in an “alternative conversion list” by their station addresses (MAC addresses).

On the initial global bridging parameters configuration screen, make sure Hop Count Reduction is left at the default No, when using TRNSB.

Encapsulation Filters

Filters enable the bridge to either selectively relay or drop a particular frame on the basis of header fields used with each of the four encapsulation methods supported by the bridging service. These encapsulation methods are as follows:

- Ethernet
- IEEE 802.2 logical link control
- IEEE 802.2 LLC with SNAP header
- Novell proprietary

Figures 33 through 36 illustrate each method of encapsulation.

Preamble 8 octets	Destination 6 octets	Source 6 octets	Type 2 octets	Data 46–1500 octets

Figure 33. Ethernet Encapsulation

Ethernet version 2 encapsulation (shown in figure 33) prefixes an eight-octet preamble, six octets of destination-address information, six octets of source-address information, and two octets of protocol-type information to the frame. It appends a four-octet frame check sequence to the frame.

DSAP 1 octet	SSAP 1 octet	Control 1 octet	Data 46–1500 octets

Figure 34. 802.2 Encapsulation

802.2 encapsulation (shown in figure 34) prefixes one octet of destination service access point identification, one octet of source service-access-point identification, and one octet of control information to the frame. The 802.2 frame, in turn, will be encapsulated within a MAC-level media-specific frame.

DSAP 1 octet	SSAP 1 octet	Control 1 octet	Organization 3 octets	Prot. Type 2 octets	Data

Figure 35. SNAP Encapsulation

SNAP encapsulation (shown in figure 35) is an extension of 802.2 encapsulation. It prefixes one octet of DSAP information, one octet of SSAP information, one octet of control information, three octets of organizational information, and two octets of upper-level protocol type information (sometimes called Ethernet type) to the frame. The SNAP structure is further encapsulated within a MAC-level medium-specific 802.x frame.

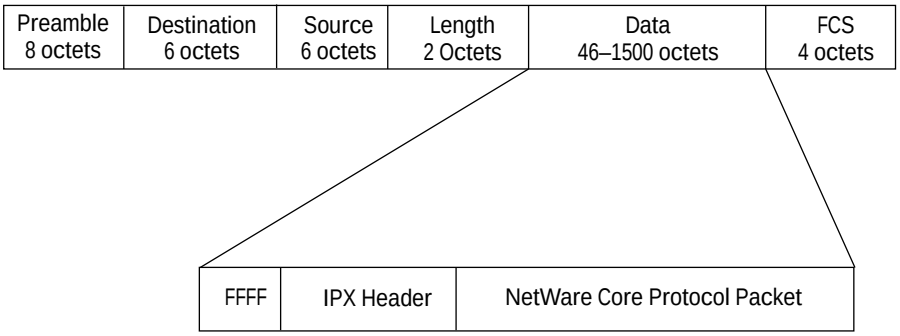

Figure 36. Novell Proprietary Encapsulation

Novell proprietary encapsulation (shown in figure 36) prefixes an eight-octet preamble, six octets of destination-address information, six octets of source-address information, and two octets of frame-length information to the unchecksummed IPX frame (indicated by a value of FFFF). It appends a four-octet frame check sequence to the frame.

Table 1 shows encapsulation support for each physical access medium:

Table 1. Encapsulation/Media Matrix

Medium	Encapsulation Method			
	Ethernet	**802.2**	**SNAP**	**Novell**
Ethernet/802.3	Yes	Yes	Yes	Yes
Token Ring	No	Yes	Yes	No
FDDI	Yes	Yes	Yes	Yes
Point-to-Point	Yes	Yes	Yes	Yes

The bridge provides a set of pre-defined filter fields. Table 2 lists encapsulation methods along with associated pre-defined fields.

Table 2. Predefined Filter Fields

Encapsulation Method	Predefined Fields
All	Source station addresses
	Destination station addresses
Ethernet	Type
802.2	SSAP
	DSAP
SNAP	Organization
	Ethertype

The bridge supplements basic filtering functionality by providing the ability to specify user-defined fields within each of the supported encapsulation formats. It also provides the ability to specify lists which contain a collection of value ranges to be filtered.

The Traffic Priority configuration parameters should be left set to Normal, the default, when using encapsulation filtering.

Traffic Filters

Traffic filters apply to all incoming bridge traffic across the circuit group. You can, if you wish, construct up to 31 filters for each bridging circuit group.

Conceptually a filter consists of a rule which identifies packets to be filtered, an action to take upon receipt of a frame that meets the conditions of the rule, and a precedence that identifies which action to take in the event of a frame that meets the conditions of more than one rule. A filter rule consists of three entities:

- A specified field (or fields) in the frame header
- A value (or range of values) associated with the field
- An operator which specifies the relationship between field and value.

A filter operator may take one of three values: ignore, match, or don't match. A filter precedence is designated by a decimal value from 1 to 31; the higher the value, the greater the precedence.

Traffic Prioritization

Prioritizing Bridged Packet Traffic

Router traffic, both bridged and routed, generally moves on a "first-in, first-out" basis. Prioritization can help to ensure that bridged packets that are sensitive to long response times (such as SNA packets) will not be delayed or dropped due to delays caused by traffic congestion. Prioritizing is based on the bridging circuit group and is done on inbound packets that will be bridged over a WAN circuit. There are two methods of prioritizing bridged packets:

- Configure the same priority level for all bridged packets within a circuit group. This option globally prioritizes all incoming bridged packets for the specified group to be processed with high, normal (the default), or low priority. This option gives you the following capabilities:
  - Configure bridged traffic to a lower priority than routed traffic (which is always normal).
  - Configure bridged traffic to a higher priority than routed traffic (which is always normal).
- Configure different priority levels for different types of bridged packets within a circuit group. This is done with traffic filters, by setting the "Action" to high or low priority rather than to accept or drop. Within a specified circuit group, this option prioritizes individual encapsulation types of incoming packets to be bridged at different priority levels.

Aligning Circuit Bandwidth to Prioritization Needs

Together with establishing priorities for bridged packets, you can also specify bandwidth allocation within a WAN circuit for each priority level. This reserves a mimimum bandwidth if higher-priority traffic would take it all.

Internet Protocol Routing Service

Product Note

Routing consists of sending a packet from a source to a destination over one of several available paths. Unlike bridges, which must store routes to all hosts (end nodes) in an extended network, routers need only store routes to other networks, and to the end nodes in directly connected networks. In the Internet Protocol (IP) environment, the packet may be referred to as a datagram.

Routers create a network of networks. Local area networks that use the Internet TCP/IP protocol suite can communicate with each other through IP routers. A TCP/IP network, or IP network, is typically an entity-wide, geographically dispersed network consisting of several subnetworks connected by routers. Example entities include companies, academic institutions, government agencies, or subgroups of these. IP subnetworks may consist of LAN or WAN networks connected to the router interfaces. A single subnetwork may consist of several LAN and/or WAN segments connected by bridges or repeaters (hubs).

See the specifications for IP listed at the end of this note.

Applications of IP

IP is the most widely implemented networking protocol, available on over 200 computer platforms. It supports the broadest set of application-level services, some of which are listed below.

- File transfer and distributed file systems:
  - File Transfer Protocol (FTP)
  - Trivial File Transfer Protocol (TFTP)
  - Network File System (NFS)
  - Network Basic Input/Output System(NetBIOS)
- Electronic mail and news:
  - Simple Mail Transfer Protocol (SMTP)
  - Privacy Enhanced Mail (PEM)
  - Remote terminal emulation
  - Telnet virtual terminal
  - rlogin
  - tn3270 access to IBM VM/CMS computers
- Window systems:
  - X Window System
- Time synchronization:
  - Time protocol
  - Network time protocol (NTP)
- Security systems:
  - Kerberos authentication and authorization system
- Network management:
  - Simple Network Management Protocol (SNMP)
  - ISO CMIP/CMIS over TCP/IP (CMOT)

An HP router will route traffic generated by these application-layer services throughout an internetwork.

IP Addressing Scheme

If your IP network will be connected with other IP networks worldwide, you must use assigned IP addresses. Otherwise, you can build your own IP addressing scheme.

Assigned Addresses

Hewlett-Packard strongly recommends that if you intend to integrate your network with other IP networks or to expand your network in the future, you use assigned addresses. There is a formal process to obtain assigned unique IP addresses for networks worldwide. Contact Government Systems, Incorporated (GSI)—formerly known as the DDN Network Information Center, or NIC—by phone, mail, or electronic mail as shown below. They will provide instructions and the necessary documents to assign and register your IP addresses.

Telephone	in U.S. only:	800-365-3642
	worldwide:	703-802-4535
Mailing Address	Government Systems, Inc. Attn: Network Information Center 14200 Park Meadow Drive Suite 200 Chantilly, VA 22021	
E-mail	hostmaster@nic.ddn.mil	

If your network is isolated and will not ever be connected to any other IP networks, you can build your own IP addresses. If you use your own addressing scheme, be aware that any connection to another IP network could cause communication problems on both networks. The addressing scheme on the two networks must be compatible, and *each address must be unique.*

For an description of IP addressing, see the book, *Interconnecting With TCP/IP: Principles, Protocols, and Architecture* (Volume I), Second Edition, by Douglas E. Comer, published by Prentice-Hall, Inc., 1991.

Description

An IP address consists of 32 bits divided into two or three fields: either network number and host number or else network number, subnet number, and host number. (An IP network generally comprises a single company or location.) The interconnection of IP networks is an internetwork. The most widely used internetwork is the Internet, which includes public and private IP networks. At the destination network, the subnet number (if any) is used to send the packet to the correct subnetwork. After that, the packet is sent to the correct host number. (A host is a node or network device that supports IP communication on the network.)

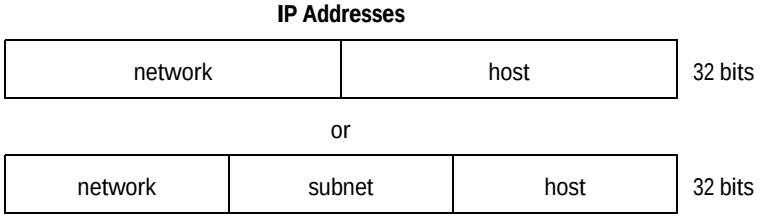

The address field length of an IP address depends on the address class.

The class A address assigns 8 bits to the network field and a total of 24 bits to the rest of the address. This address class can address almost 17 million different nodes on an IP network. The class A address is used for very large networks.

The class B address assigns 16 bits to the network field and a total of 16 bits to the rest of the address. This can address over 65 thousand different nodes on a network.

The class C address assigns 24 bits to the network field and 8 bits to the rest of the address. This can address 254 different nodes on a network.

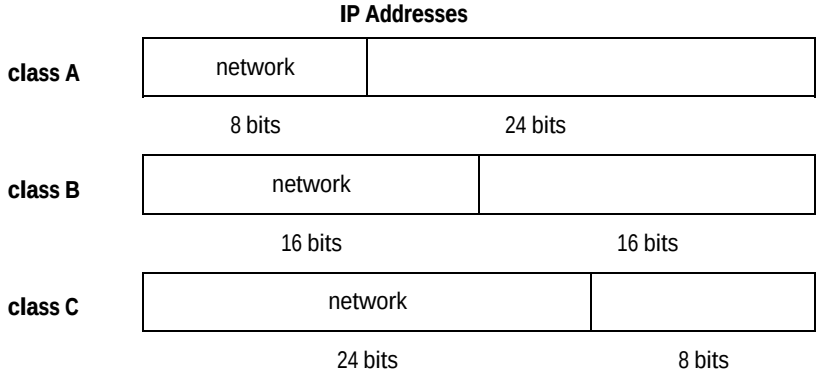

After you have selected the address class and network number, the rest of the address bits are allocated to the host field or subdivided into subnet and host fields. The field lengths chosen for these fields will depend on how the network is subdivided. For example, a class B address could have 8 bits assigned to the subnet field and 8 bits to the host field. This would allow up to 254 subnetworks with up to 254 hosts on each. Or, it could have 9 subnet number bits and 7 host number bits. This would allow about 500 subnetworks with up to 126 hosts on each.

Dividing the network into subnetworks helps with the administration of a large network. When selecting the lengths for the subnetwork and host fields, consider the total number of subnetworks that will potentially be needed on the network and the total number of nodes in each subnetwork. Make sure that the lengths of each field will accommodate these projected totals and leave room for expansion. (See “Suggestion”, following the discussion of subnet masks below.)

The actual subnet and host number values are assigned by the network administrator. Neither the subnet or host number can be all 1s or all 0s (255 or 0). These are reserved addresses. You might want to consider using a network number assigned by the Network Information Center at Government Systems, Inc., even if you are not currently connecting to the Internet. If you do this now, you will not have to reconfigure your IP network if you connect to the Internet later.

IP Addresses

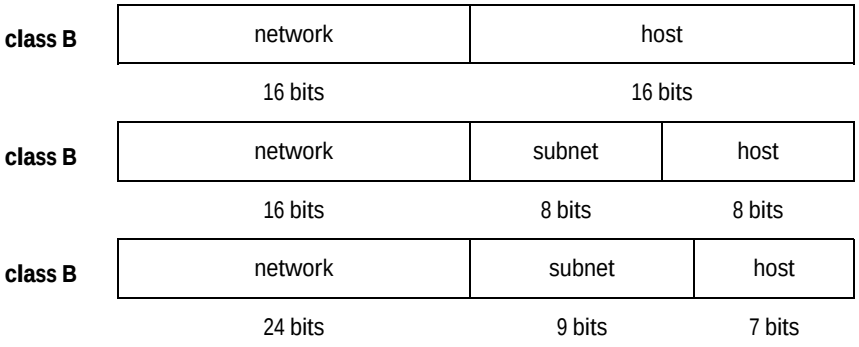

Notation

IP addresses are written in dotted decimal notation. Each decimal group (between the decimal points) is the decimal equivalent of 8 bits of the binary address. Notice that the dotted decimal divisions do not exactly correspond to the network, subnet, and host field divisions of the address. One address field may cover more than one dotted decimal division. Or, the division between address fields may not fall in the same place as a dotted decimal division. It is recommended that subnet/host fields be allocated at a dotted decimal division point (at a "byte boundary") whenever possible.

IP Addresses

binary:	1000 0100	0000 0111	0011 0100	0001 0011
dotted decimal:	132.	7.	52.	19
address fields:	network		subnet	host
	132.7		52	19

The address class can be determined by the first decimal number. This is because the leftmost bits are assigned according to the address class:

For a class A address, the first digits are in the range 1 through 126 (leftmost bit is 0).

For a class B address, the first digits are in the range 128 through 191 (leftmost bits are 10).

For a class C address, the first digits are in the range 192 through 223 (leftmost bits are 110).

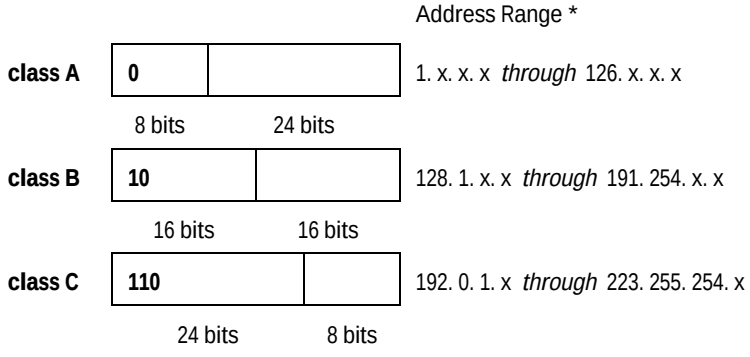

* Some addresses are reserved and are not included in the ranges listed.

Subnet Mask

When assigning IP addresses, you will also assign subnet masks. A subnet mask tells you the total length chosen for the network and subnet fields. It is constructed as follows:

1. "1" is assigned to each network and subnet bit.
2. "0" is assigned to each host bit.
3. Each group of 8 bits is converted to its decimal equivalent to obtain dotted decimal notation.

For example, the subnet mask for an IP address with field lengths of network=16, subnet=8, and host=8 is 255.255.255.0. The subnet mask for an IP address with field lengths of network=8, subnet=8, and host=16 is 255.255.0.0.

IP Addresses

binary:	1111 1111	1111 1111	1111 1111	0000 0000
subnet mask:	255.	255.	255.	0
address fields:	network		subnet	host
	16 bits		8 bits	8 bits

Notice that the field division may not correspond neatly to the dotted decimal divisions as in the above examples. For example, the subnet mask for an IP address with field lengths of network=16, subnet=7, and host=9 is 255.255.254.0.

IP Addresses

binary:	1111 1111	1111 1111	1111 111 0	0000 0000
subnet mask:	255.	255.	254.	0
address fields:	network		subnet	host
	16 bits		7 bits	9 bits

Suggestion for Assigning Addresses

Once your network number is assigned and you have selected the subnet mask, you have apportioned the address space that will be available for additional subnets and for additional hosts in the future. Because it may not be clear which will increase more—subnets or hosts—you can start with the scheme described below to reserve the most flexibility for expansion. You want to reserve the option of slightly changing your subnet mask in the future (to adjust the relative allocation of space for new subnets and for new hosts), without having to also change the IP address configured on every device in your network!

Once you have established your subnet mask, start assigning subnet numbers at the most significant bit of the subnet field, and host numbers at the least significant bit of the host field. For example, starting with a class B address and subnet mask 255.255.255.0, start assigning subnets 128, 64, 192, 32, 160, etc. Within each subnet, start assigning nodes 1, 2, 3, 4, 5, etc. Figure 1 illustrates an example class B network.

First Subnets		First Hosts	
Binary	**Decimal**	**Binary**	**Decimal**
1000 0000	128	0000 0001	1
0100 0000	64	0000 0010	2
1100 0000	192	0000 0011	3
0010 0000	32	0000 0100	4
1010 0000	160	0000 0101	5
0110 0000	96	0000 0110	6
1110 0000	224	0000 0111	7

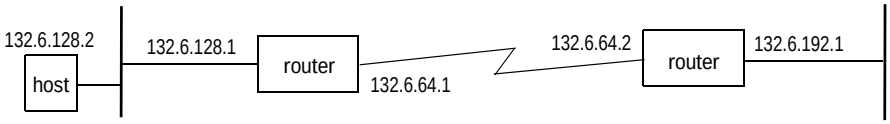

Figure 1. Initial IP Address Asssignments in Router Network 132.6

After it's clear whether more subnets or more host numbers within each subnet are used, you may be able to change your subnet mask from 255.255.255.0 (using 8 bits for subnets) to 255.255.255.128 (using 9 bits for subnets) or to 255.255.254.0 (using 7 bits for subnets).

Example Topology

Figure 2 shows an IP internetwork that connects subnetworks in five cities using HP routers, with IP routing service enabled, and an HP Remote Bridge RB. The network is an autonomous system with IP network address 128.1.0.0 and subnet mask 255.255.255.0. The 16 most significant bits (128.1) identify the network. Subnetworks are identified in the figure by letters A through K. WAN links from New York to Atlanta and Paris are each unique subnetworks E and H, respectively. The New Orleans bridged LAN and WAN links form a single subnetwork A. The link between Paris and Lyon through the Transpac X.25 packet-switching network forms a single subnetwork J.

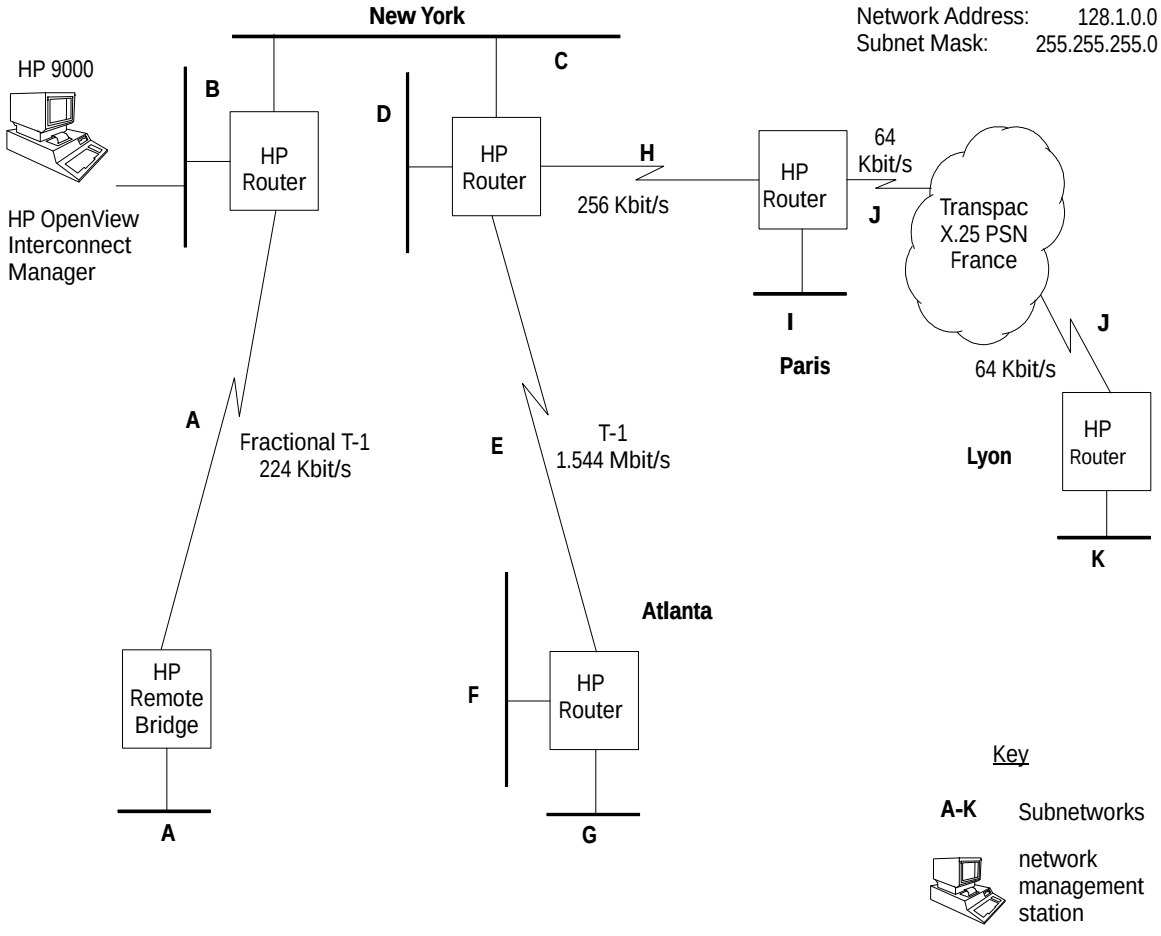

Figure 2. Example IP Network

IP Routing Decisions

IP routing decisions are based upon the destination network-layer address contained in each data packet that is traveling through the IP network. The most significant bits of the address identify the destination subnetwork, while the least significant bits identify a specific node (host, router, HP managed hub, HP managed bridge, etc.) on that subnetwork. The subnet mask is used to distinguish these parts of the 32-bit IP address.

In figure 2, the subnet mask 255.255.255.0 determines that the most significant 24 bits of each IP address will identify the subnetwork, leaving 8 bits on each subnetwork to identify individual hosts and router interfaces. All addresses will be of the form 128.1.*sss.hhh*, where "*sss*" is a number (1–254) identifying the subnetwork and "*hhh*" is a number (1254) identifying a node (host) on that subnetwork. Note that each port (network interface) on the router resides on a different subnetwork and requires a unique IP address. For more information on IP addressing, refer to the student workbook in the HP course "*HP Internet Routing Products*", or to the reference manual for the HP routers.

The router bases IP packet-routing decisions on the following:

- The network and subnetwork portion of the destination IP address contained in each packet.
- The information it has about the network's topology. This topology information is maintained in the "IP routing table".

IP Routing Table

The IP routing table contains an entry for each subnetwork that the router has learned about. For each destination subnetwork entry, the table also contains the following information:

- *Metric:* the cost, typically in terms of "hop count", to the destination subnetwork from this router. A hop count of "*n*" indicates that *n* routers separate this router and the destination subnetwork.
- *Next Hop:* If the destination subnetwork is directly connected to this router, then the next hop is the IP address of this router's interface to the destination subnetwork. If the destination subnetwork is a remote subnetwork, the next hop is the IP address of the inbound interface of the next router along the path to the destination.
- *Route Type* specifies whether the route is remote or directly connected. Also, if the router is notified that a learned route is no longer available, then the route is marked invalid. The route will remain in the table until the router is notified of a valid route or is rebooted.
- *Route Learned* indicates whether the route was statically learned (that is, the route was specifically configured) or learned dynamically through one of the routing protocols (RIP, OSPF, or EGP).
- *Age* indicates the number of seconds since the route was learned.
- *Interface:* a number, assigned by the router, indicating the port through which the next hop is reached.

For more information on the IP routing table, refer to the reference manual for the HP routers.

Note that while the IP protocol uses the information in the routing table to determine where to deliver each packet, the IP protocol does not gather the information to create and maintain the routing table. That function is performed by one or more of the dynamic routing protocols: RIP, OSPF, or EGP. Alternatively, routes can be statically configured.

Routing Protocols

RIP Routing Information Protocol, RIP, is an interior gateway protocol (IGP) for exchanging network reachability and routing information within an autonomous system. RIP is relatively simple to configure and is best suited for smaller networks (fewer than 15 hops in diameter), although it can be configured for a network diameter of up to 127 hops.

Note that using a larger network diameter makes the network prone to the "slow convergence" or "count to infinity" problems. This is a state in which some routers have inconsistent information. Also, the network may have routing loops because routing advertisements must propagate further across the network. To avoid the slow convergence problem, the HP routers support RIP with "split horizon updates" and "poison reverse with triggered updates". To enable RIP to account for different link speeds, the HP routers allow you to assign a higher cost to interfaces with lower-speed networks. For further explanation, refer to the "IP Network Interface Definition" section in the IP chapter in the reference manual.

OSPF Open Shortest Path First, OSPF, is an interior gateway protocol (IGP) for exchanging network reachability and routing information either within or between autonomous systems. Configuring OSPF is more involved than configuring RIP. However, it is much more robust and is better suited to larger, more complex networks with widely varying link speeds.

To reduce the amount of traffic generated by advertisements between routers, OSPF supports "areas" and "designated routers". A network can be broken into several areas, each consisting of a group of subnetworks and routers. Within each area, a designated router will be elected. Other routers in the area will exchange link-state advertisements (LSAs) only with the designated router, not with each other. Additionally, only those routers with direct connections to multiple areas will exchange advertisements about the subnetworks within each of those areas.

OSPF enables you to configure variable-length subnet masks on different subnetworks, which can be used to conserve IP address space. OSPF enables you to configure a password so that all OSPF messages received will be authenticated for added network security. Like RIP, OSPF allows you to configure interface cost, but has a maximum value of 65535. Additionally, OSPF enables you to configure the intervals for the transmission of the following:

- Messages to elect the designated router
- Messages to determine whether a silent router is down
- OSPF advertisements

For more information on RIP and OSPF, refer to the reference manual for the HP routers.

EGP Exterior Gateway Protocol, EGP, is a protocol for exchanging network reachability information between routers in different autonomous systems. When connecting an autonomous system to the Internet, every autonomous system must use EGP to advertise network reachability to the Internet's core gateway system. The Internet is a globally administered confederation of autonomous systems (networks).

Note that any or all of the supported routing protocols may be active on the router simultaneously. Each of the routing protocols derives its routes using different algorithms, presenting the possibility of having multiple routes for a single destination. The routing table, however, can contain only one entry for each destination subnetwork and will select that entry with the following order of precedence:

Highest Preference:	Static Routes
	Routes derived from OSPF
	Routes derived from EGP
Lowest Preference:	Routes derived from RIP

For more information on EGP and static routes, refer to the IP chapter in the reference manual.

The IP Network Interface Definition

The HP router must have an IP interface defined for each attached network using IP routing. On the HP routers, “circuit groups” connect the router to its networks. On the other end of each circuit group on this router is, in the case of a WAN link, another router or gateway or host, or in the case of a LAN link, a network or subnetwork.

If one link has more than one network or subnetwork attached, then each subnetwork must have its own network interface definition with its own unique IP address. Note that the nodes on the two subnetworks will always communicate with each other through the router even if they are on the same cable. For more than one subnetwork on a port, configure that circuit group in more than one network interface definition, each with a different IP address. Suppose, for example, that LAN interface B, shown in figure 3, accesses networks N_x and N_y. In this case, you need two IP addresses. This is because one address for interface B is needed to access LAN x, and another address for interface B is needed to access LAN y.

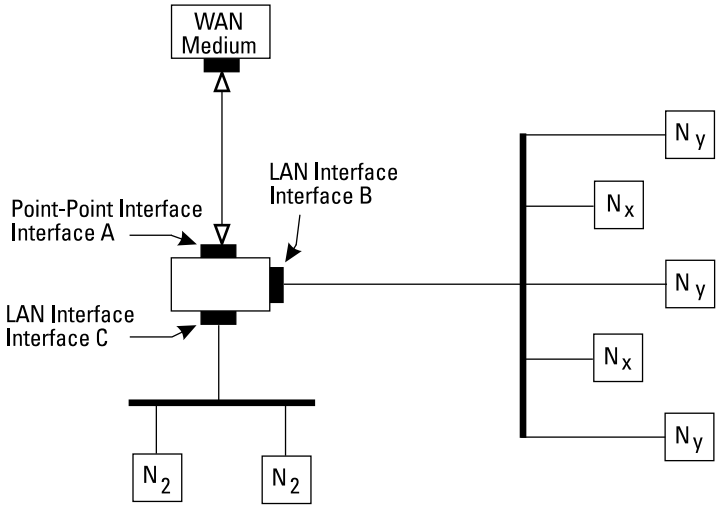

Figure 3. Sample Topology for Multiple Network Interface Definitions

Static Routing

The IP router provides the following types of static routing that can be used instead of dynamic routing, based on a packet's destination address:

- Conditional and nonconditional static routes to specify a path to another router for a specific destination.
- Default routes to specify a path to another router for all routes not explicitly known.
- Adjacent host routes to specify a path to a host (end node).

Static Routing Examples

Some cases in which these types would be used are the following:

Static route example 1: To reach a network through another router that does not support RIP, a static route must be configured, since this HP router will not receive routing information about the other network. See figure 4 for an example. For a node on the Red net to reach a node on the Green net, on this router configure a static route on network interface 1, using the destination IP address of router X's interface B, and using the IP address of router X's interface A as the next hop. The cost is 1. Router X also must have a static route configured to reach the Red net.

Adjacent host route example: To accommodate individual hosts that do not implement the ARP or HP Probe protocol, configure an adjacent host route to each such host. The host must be on a LAN directly attached to this router. You will configure the IP address and station address of the host on the attached LAN, and the subnet mask of the host's subnet.

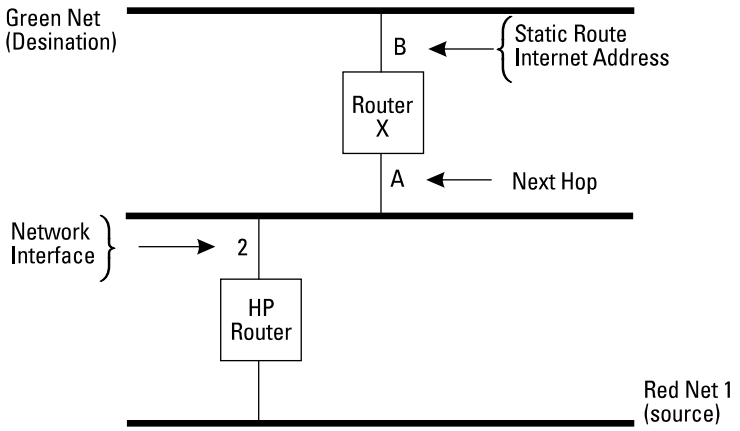

Figure 4. Static Route Example 1

Default route example 1: When the HP router has a small number of directly connected networks and has a single connection to another router or router backbone, a default route can be configured to the other router. The example shown in figure 5 shows that routing updates need not be sent from router 2 to router 1, because all traffic not for the Red or White networks can be sent to router 2, thereby preserving bandwidth inbound to router 1. The Red and White networks attached to router 1 are included in routing tables on all the other routers, so the Red and White networks are reachable by all. Network interface 3 on router 1 is a default route to the other routers; interface A on router 2 is the next hop. RIP Supply will be configured, but RIP Listen, Default Route Supply, and Default Route Listen will not be configured. Cost is not applicable. Network interface A on router 2 is not configured as a default route. RIP Listen will be configured, but RIP Supply, Default Route Supply, and Default Route Listen will not be configured.

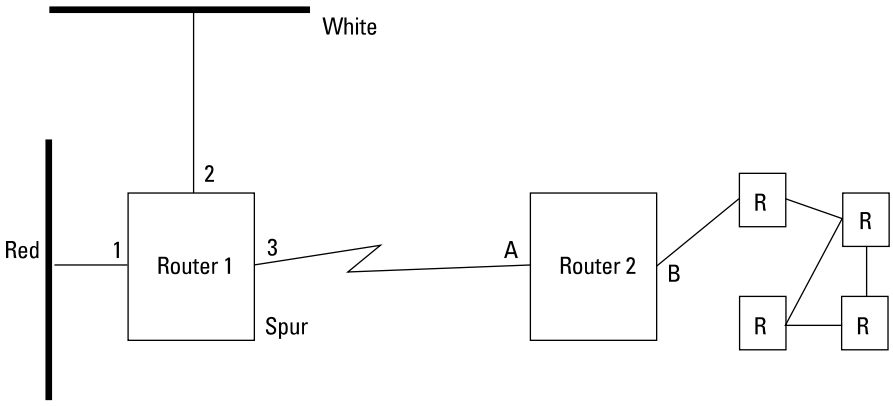

Figure 5. Default Route Example 1

Default route example 2: To connect an HP router to another network that uses a different interior gateway protocol (IGP), a default route is used. One such example, shown in figure 6, is connecting an HP router to a Cisco router backbone that uses IGRP, a proprietary routing protocol. Network interface 3 on the HP router is a default route to router X; interface A on router X is the next hop. RIP Supply, RIP Listen, Default Route Supply, and Default Route Listen will not be configured. Cost is not applicable. Also configure static routes on router X to the Red network and to router 2.

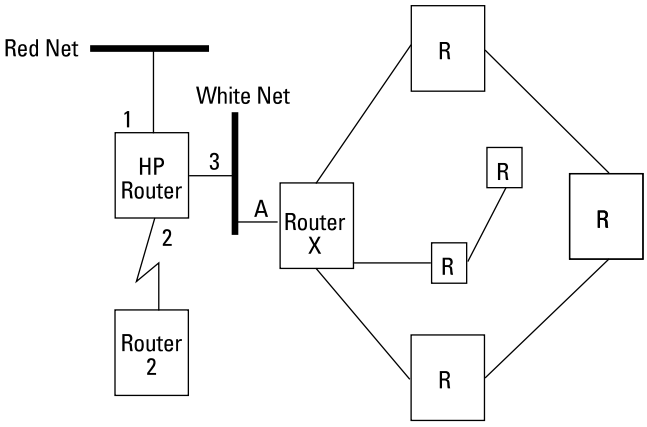

Figure 6. Default Route Example 2

Static route example 2: To restrict packets, for security reasons, to paths you specifically configure, all routes would be configured statically. Note that one of your redundant links could never be used for

IP traffic. RIP Supply, RIP Listen, Default Route Supply, and Default Route Listen will not be configured. In addition, for LANs and direct point-to-point connections, ARP or both ARP and HP Probe will be used for address resolution; Normal ARP should be on.

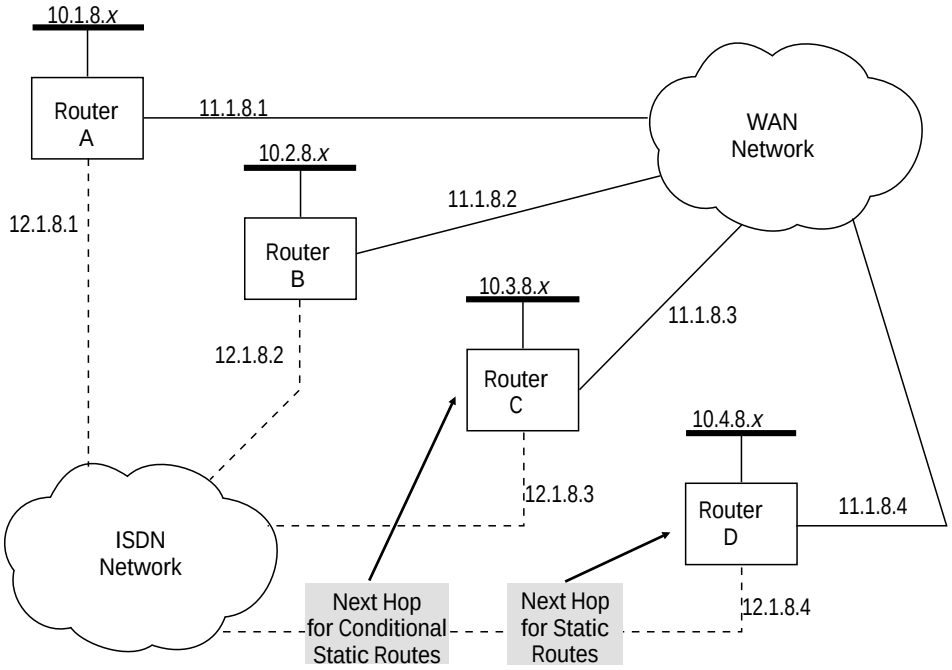

Figure 7. Conditional Next Hop Routing

Conditional Static Route

A static route can be configured as conditional on the status of a circuit group other than the one that directly accesses the next hop router. For example, consider figure 7.

- The frame relay subnet 11.1.8.*x* provides the primary routes.
- The ISDN (V.25 bis) subnet 12.1.8.*x* provides the backup for the frame relay subnet.
- RIP is run over the WAN network, enabling all routers to learn about the 10.*x*.*x*.*x* networks attached to other routers (when all WAN links are up).
- Routers A, B, and C have static routes (using the ISDN or dial-up network) to all the non-locally-attached 10.*x*.*x*.*x* networks. The routing cost for these static routes is configured to be higher than for the WAN routes, which means that a static route to a particular router will not be used unless the WAN/RIP route to that router goes down.
- On the ISDN/dial-up static routes to the 10.*x*.*x*.*x* networks, the next-hop router for routers A, B, and C is router D. Router D has ISDN/dial-up static routes to all of the 10.*x*.*x*.*x* networks, but the next-hop for these static routes is the router that is directly connected to each particular 10.*x*.*x*.*x* network. For example, the next hop for router D to network 10.1.8.*x* through the ISDN/dial-up static route is 12.1.8.1.

The configuration in figure 7 solves any single-point WAN failure to routers A, B, or C. For example, if the WAN link to router B breaks, traffic between B and all other routers will be conducted through the ISDN/dial-up static routes to D. (D will advertise the new route to B over its frame relay connection.) If, however, the WAN line to D breaks, D will either constantly be calling A, B, or C, or they will be trying to call D when they have data for the 10.4.8.*x* network accessed through D. Because there would not be enough ports on D to handle all of the calls that could occur simultaneously, the above layout is inadequate.

The solution is to implement conditional static routing that will replace the existing static routing in the event of a break in the WAN connection to router D. That is, if the circuit group on the WAN route to D breaks, the existing (ISDN) static routes connecting D to A, B, and C are bypassed for a new set of static routes configured on D, with C as the next-hop router. The result is to shift the next-hop router status from D, which has only one static route for handling all incoming and outgoing traffic, to C, which has both a WAN and an ISDN route available.

To achieve this effect, router C will be configured to advertise its ISDN route to D as lower in cost than the individual ISDN static routes from A and B to D. (Note that this ISDN route from C to D would be configured to be more costly than the WAN route when the WAN route to D is up.) The following table shows an example of how to configure the above solution for access to 10.4.8.*x* with and without a break in the WAN route to router D.

Route Type	Internet Address (Destination)	Next Hop Router	Condition	Cost to Destination
Primary	10.1.8.*x*	11.1.8.1 (A)	n/a	1
	10.2.8.*x*	11.1.8.2 (B)		
	10.3.8.*x*	11.1.8.3 (C)		
Conditional Static	10.1.8.*x*	12.1.8.3 (C)	Use if 11.1.8.4 (WAN) goes down	3
	10.2.8.*x*			
	10.3.8.*x*			
Static Only	10.1.8.*x*	12.1.8.1 (A)	n/a	5
	10.2.8.*x*	12.1.8.2 (B)	n/a	5
	10.3.8.*x*	12.1.8.3 (C)	n/a	3

Preference and the Routing Pool

The IP router maintains a routing pool which contains information supplied by up to three routing protocols (RIP, EGP, OSPF) in addition to statically configured routes. Consequently, the routing pool may contain multiple routes to the same destination. Each route carries an associated preference value that determines the “best” route where more than one is available. The forwarding table is constructed with this best route for each known destination. The routing pool is updated in response to received protocol traffic; updates are subsequently reflected in the forwarding table. The pool also provides the database used by the routing protocols to prepare their link state/routing advertisements. You can mediate the flow of routing data to and from the routing pool; see “Routing Filters”, the next section after “Static Routing”.

By default, the IP router uses manually configured static and/or default routes in preference to routes gathered by protocol exchanges. You can configure the preference for each static (conditional or nonconditional) route as a weighted value used by the IP router to select from multiple routes to a single destination. 0 is lowest preference; 16 is highest preference. Routes with higher values will be selected for IP routing in preference to routes with lower values.

Default Route

A default route is the path that a router directs a packet to when its routing table does not contain the destination network specified in the packet's IP header.

Upon receiving a packet, a router first compares the packet's destination network against those found in its internal routing table. If the router finds the destination network in its table, it directs the packet toward the corresponding interface. If it doesn't, it directs the packet toward the interface associated with the default route (providing one is configured) for further processing by a neighboring router.

You can configure up to four default routes. By configuring multiple routes, you ensure that a packet can be re-routed if the interface associated with the default route is disabled.

You assign preference values (between 1 and 16, with 16 being the highest) to each default route, to determine which default route gets the highest priority. If the highest priority default route is unavailable, the router uses the next most-preferred default route. If this interface is also unavailable, the router then chooses the next most-preferred default route. Should a disabled default route with a higher priority value re-enable, the router uses it as the default route.

For example, the figure 8 shows a sample multiple default route topology and the preference values assigned to the four default routes. The IP router directs packets addressed to networks A, B, or C through interfaces A, B, and C. All other packets are directed to the default route D because it has been assigned the highest preference among the default routes available. If interface D becomes disabled, then the router uses default route E (and so forth for default routes F and G).

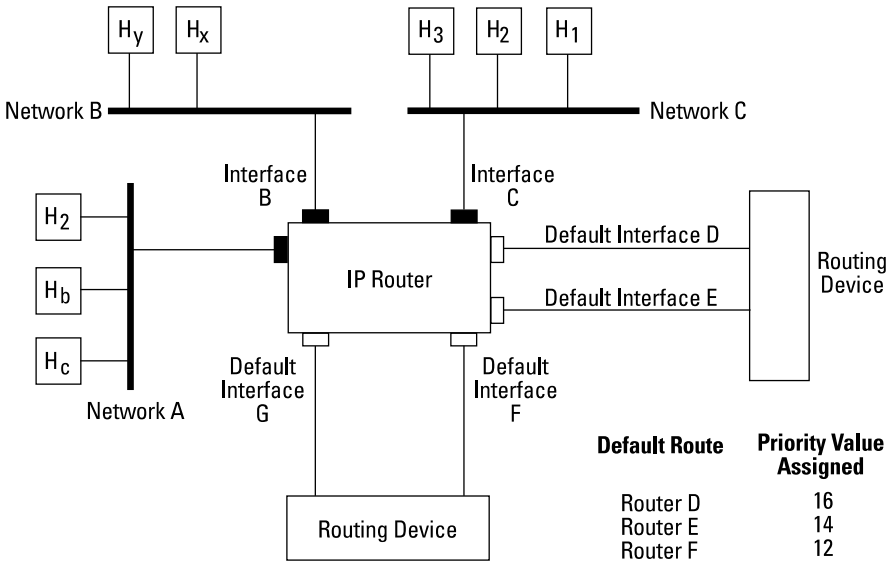

Figure 8. Multiple Default Routes

Adjacent Host Route

Adjacent hosts are end nodes on a locally-attached network. Specify an adjacent host if you are setting up a network or if a particular local host or hosts don't respond to ARP requests.

The static routing type is configured to be adjacent host, and one of three encapsulation methods is configured, as follows.

Ethernet (the default) is the standard Ethernet 2.0 encapsulation for hosts that support Ethernet. This type is required for point-to-point or any type of X.25 interface.) If you are defining a LAN interface (Ethernet or IEEE 802.*x*), you must specify the encapsulation method supported by the attached network.

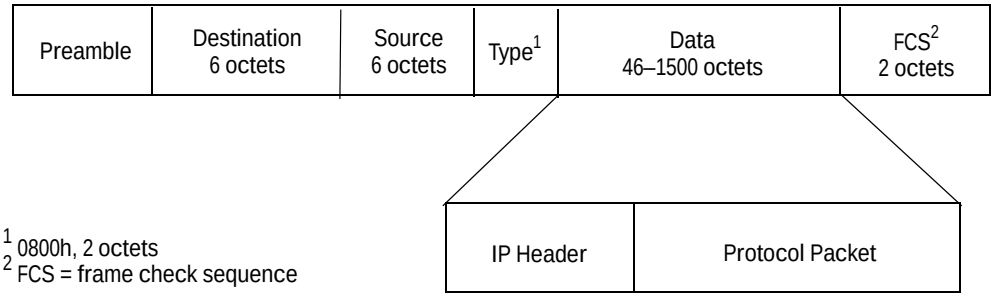

Figure 9. Ethernet Encapsulation

802.2 can be used for hosts supporting IEEE 802.2 over IEEE 802.3 LAN interfaces. The 802.2 structure is encapsulated as shown, and is further encapsulated within a medium-specific 802.*x* packet.

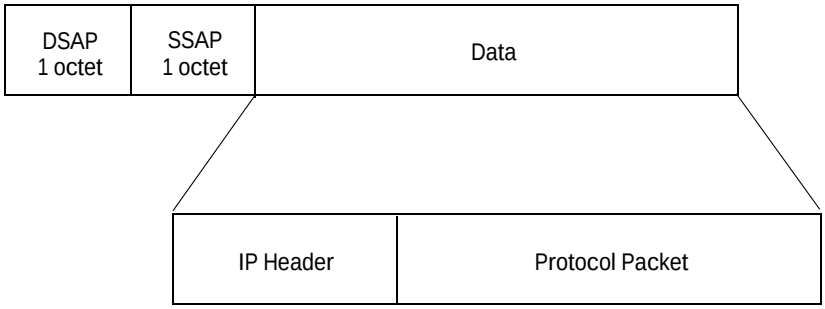

Figure 10. IEEE 802.2 Encapsulation

SNAP (an extension of 802.2 encapsulation) can be used for hosts that support SNAP. The SNAP structure is encapsulated within a medium-specific 802.*x* packet.

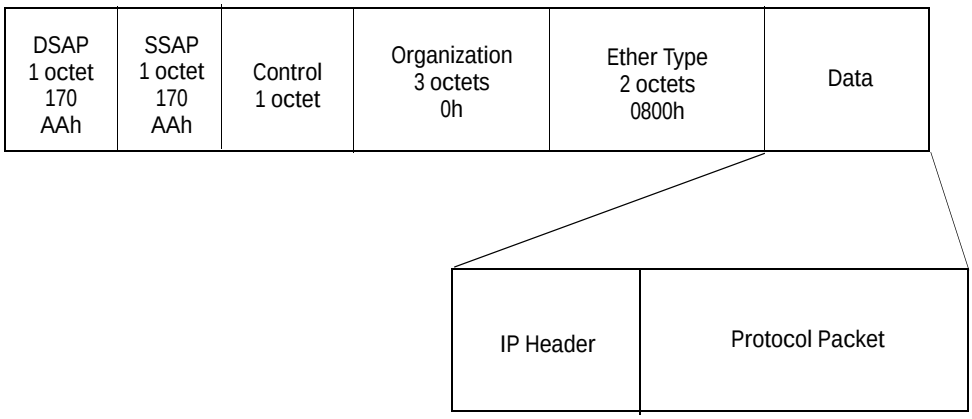

Figure 11. SNAP Encapsulation

IP Filters

Routing Filters

The HP routers support import and export route filters to allow you to modify the order of precedence for deriving routes to a given destination network. Import filters restrict the input of routing information about a given destination by a particular routing protocol. Export filters restrict the output of routing information about a given destination by a particular routing protocol. (Refer to the description of the routing pool, “Preference and the Routing Pool”, within the preceding section on static routing.)

RIP import rules govern the addition of new routes to the routing pool. RIP maintains a distinct set of import rules. For example, upon receiving a new routing update, RIP consults its specific import rules to validate the information before inserting the update in the routing pool. Import rules contain search information (used to match fields in incoming routing updates) and action information (used to specify the action to take with matched fields).

RIP export rules govern the propagation of routing information by RIP. RIP maintains a distinct set of export rules. For example, when preparing a routing advertisement, RIP consults its specific export rules to determine whether routes to specific networks are to be advertised and how they are to be propagated. Export rules contain network numbers (used to associate a rule with a specific network) and action information (to specify a route propagation procedure).

The relationship between the routing pool, forwarding tables, and the import and export rules is shown conceptually below.

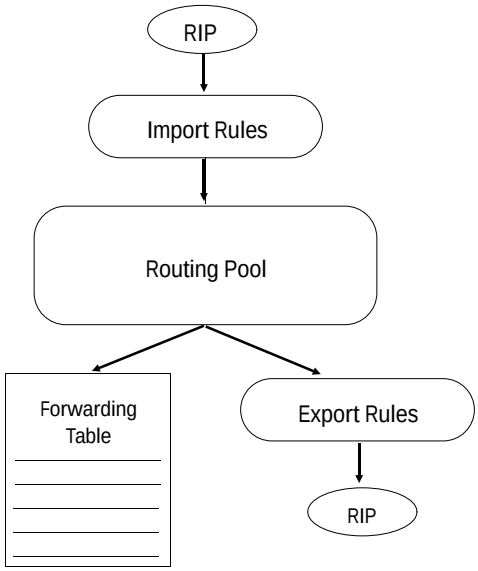

Figure 12. Routing Information Data Flow for RIP

Constructing RIP Import and Export Route Filters

- Each filtering rule must specify an incoming/originating routing protocol. Rules in the form “Accept Network Number Mask” are invalid.
- The routing pool holds one route per routing protocol per network. If there are multiple import rules for a specific network number and protocol, the latest routing update received using any of these rules is the route active in the routing pool.
- Most rules include a network number and mask; however, the network number and mask are not always required.
- Each import rule can contain up to three search fields; however, it need not contain any. As a result “Ignore RIP” is a valid rule. If such a rule were not specified, any received RIP packets that did not match a rule would be accepted into the routing pool.
- If a rule is created and contains a network address that is a superset of an address from another rule, a packet matching the more specific address is not affected by the rule for the less specific address. The only exception is a rule that specifies no network address.

Packet Filters

In addition to routing protocol filters, HP routers also support packet filters that can be used to secure the network or to control traffic flow. Packets can be forwarded, dropped, or passed to subsequent filters based the contents of specific fields within the IP packet, UDP packet, or TCP segment headers—examined either singly or in combination. You can filter on the source or destination IP address of the network or node. You can also filter on the application that sent the packet, using the TCP/UDP source or destination port on each network interface.

TCP/UDP port filters use well-known transport-layer port numbers for higher-layer services such as Telnet, FTP, and TFTP. The User Datagram Protocol (UDP) and Transmission Control Protocol (TCP) are internet transport-level protocols. TCP provides a reliable connection mode packet service while UDP provides connectionless datagram service. UDP packets and TCP segments are originated by ports and addressed to ports. Ports are logical abstractions used by transport-level protocols to distinguish between multiple sources and destinations at a single host.

To facilitate application-to-application data flow, the Internet has assigned well-known port numbers to certain commonly used application programs. Examples of well-known port numbers include port numbers assigned to remote-login (Telnet) programs, file-transfer programs, and remote-job-entry (RJE) programs.

TCP and UDP Well-Known Port Numbers

Port	Protocol	Usage	Port	Protocol	Usage
0	reserved	–	42	NAMESERVER	TCP & UDP
1	unassigned	–	43	NICNAME	TCP & UDP
2	unassigned	–	53	DOMAIN	TCP & UDP
3	unassigned	–	67	BOOTPS	TCP & UDP
4	unassigned	–	68	BOOTPC	TCP & UDP
5	RJE	TCP & UDP	69	TFTP	TCP & UDP
7	ECHO	TCP & UDP	75	private dial	TCP & UDP
9	DISCARD	TCP & UDP	77	private RJE	TCP & UDP
11	USERS	TCP & UDP	79	FINGER	TCP & UDP
13	DAYTIME	TCP & UDP	95	SUPDUP	TCP
15	NETSTAT	TCP & UDP	101	HOSTNAME	TCP
17	QUOTE	TCP & UDP	102	ISO-TSAP	TCP
19	CHARGEN	TCP & UDP	113	AUTH	TCP
20	FTP-DATA	TCP	117	UUCP-PATH	TCP
21	FTP	TCP	123	NTP	TCP & UDP
23	TELNET	TCP	133-159	UNASSIGNED	–
25	SMTP	TCP	160-223	UNASSIGNED	–
37	TIME	TCP & UDP	224-241	UNASSIGNED	–
39	RLP	TCP & UDP	247-255	UNASSIGNED	–

Filter Rules Filtering decisions are based on user-defined rules. An IP filter rule consists of:

- An IP/UDP/TCP field or fields
- A value or list of values (see “Filter Lists” below)
- An operator—*match or don’t match*—to specify the relationship between the contents of the field and the value
- An action—*drop or forward*
- A filter precedence

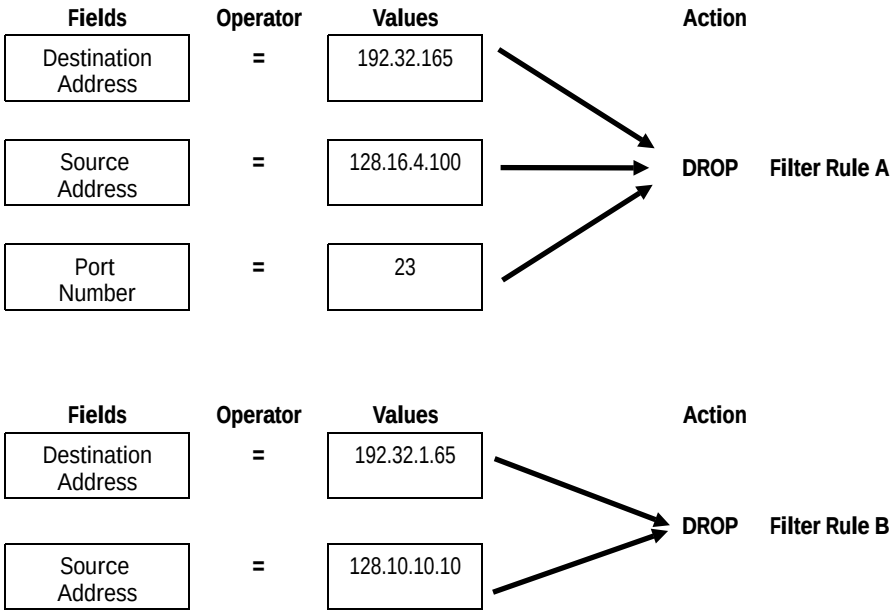

Figure 13. Sample IP Filters

IP Filter Lists Filter lists (while not required) may facilitate the configuration of filters if you wish the filter to apply to non-contiguous value ranges. A list specifies the filter values within a filter rule. The elements of a list are a name to identify the list and one or more pairs of numbers that specify a range. When a filter specifies a list name, packets are checked against the range(s) of values specified by the list. Two types of lists are defined for the two types of packet filters:

- IP address lists specify ranges of IP network addresses. Specify the lower and upper boundary of each IP address range to be in the list (or just the lower address for a single value).
- IP port lists specify ranges of TCP/UDP port numbers. Specify the lower and upper boundary of each port range to be part of the list (or just the lower port number for a single value).

User-Defined Fields in the Packet You can filter IP traffic based upon specified bit patterns contained within the IP header or the header of the upper-level protocol. User-defined field filters can be used by themselves or in conjunction with IP address and/or UDP/TCP port filters. You can specify the location of the bit pattern to filter, by giving an "offset" position (starting position, offset from the beginning of the selected header type) and a length (ending position). See figure 14.

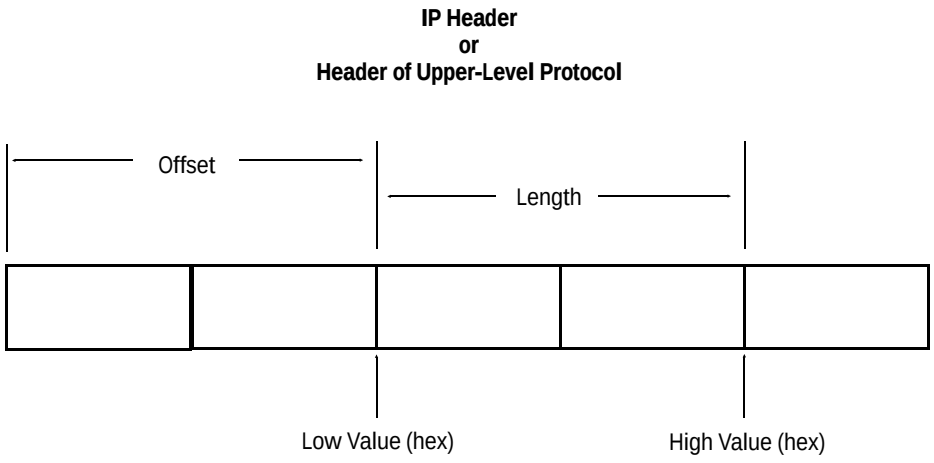

Figure 14. User-Defined Field Positioning

Device Management Functions

Ping (Packet InterNet Groper) tests reachability to IP devices using an ICMP echo request and reply sequence. ICMP, the Internet Control Message Protocol, handles IP error and control messages. Ping is used to verify and troubleshoot IP networks.

Telnet supports a remote terminal session that gives you all the capabilities you would have from a directly connected console terminal. A Telnet session can also be established from the router.

TFTP (Trivial File Transfer Protocol) is a file transfer protocol available for moving the router configuration or operating system to and from other devices on the network.

Bootp (Bootstrap Protocol) allows an HP router to be a configuration server for other HP routers, or to be a relay agent for downloading the configuration file from a file server. The configuration service on the HP router, with Quick Remote, allows branch routers to automatically connect to the network and get a minimum routing configuration from a central HP router to boot themselves.

Time Protocol allows an HP router to be a time server for other network devices, or a client of a time server, or both.

SNMP Agent (Simple Network Management Protocol agent) enables the router to be managed by network management applications across the network. HP OpenView products can access the vendor-proprietary objects on the HP routers and can create configuration files for the router. Additionally, many of the local router management operations, such as viewing the routing and address tables, are SNMP-based. The HP routers support the standard MIB-I variables, as defined in RFC 1156, as well as private-enterprise MIB variables that use the following pathname prefix:

iso.org.dod.internet.private.enterprise.wellfleet.commServer.wfmib

or

1.3.6.1.4.1.18.1.1

The standard MIB-I structure is also reproduced within the private-enterprise section (as illustrated above) at the next level under “mib”, which refers to the Internet MIB. The MIB-I variables in the private-enterprise section have different names but have the same identification codes (following the private-enterprise prefix shown above) as the router variables in the standard MIB section.

BootP and DHCP

Bootp (Bootstrap Protocol) is a protocol that runs over UDP. It uses two UDP port numbers, 67 and 68. UDP port 67 specifies a Bootp server. UDP port 68 specifies a Bootp client. In operation, a client sends a Bootrequest to a server using a destination port of 67. The server then sends a Bootreply back to the client using a destination port of 68.

DHCP (Dynamic Host Configuration Protocol) is an extension to Bootp to allow hosts to obtain more configuration parameters and temporary IP addresses from a DHCP server. DHCP is backwards compatible with Bootp, to allow a Bootp client to communicate with a DHCP server. DHCP and Bootp use the same packet formats.

An HP router can be a Bootp client and a Bootp relay agent. As a client, it can boot from a Bootp server or DHCP server. As a relay agent, it can relay both Bootp and DHCP packets. The HP router also can be a Bootp server, but only to boot another HP router that is directly attached on one of its WAN links. The HP router cannot be a file server, so it is not a Bootp server with the ability to download an entire boot file (configuration file) to the other router. In its replies to Bootp requests, it can supply the essential configuration parameters needed by an HP router to boot and *route*, and it can supply the names of the Bootp file server and full boot file. This Smartboot use of Bootp and IP is configured on the HP routers by default. See the discussion of “Instant On” and Smartboot earlier in this chapter.

In the remainder of this discussion, “Bootp” is used to refer to the combination of Bootp and the DHCP extension.

The most common networking application of Bootp is for downloading operating system code to a diskless workstation (client). The client may not know either its IP address or the IP address of its boot server. It broadcasts a Bootrequest packet to the IP broadcast address, with the client station (MAC) address. A Bootp server that receives the request looks up the client’s station address in its server table to discover what the client’s IP address and boot file name should be. The server then sends this information back to the client in a Bootreply packet addressed to the client at the MAC and IP level. The client then uses TFTP to get its boot file from the server. Another application is to download full configuration files for routers from file servers, such as from HP network and router management applications.

Bootp Relay Agent

In these applications over the network, if the client and server reside on different subnets, then the Bootrequest must be relayed by a Bootp relay agent.

The relay agent sends the Bootrequest either to specifically configured addresses or to the broadcast IP address. When the Bootp server receives the relayed request, it sends the Bootreply to the relay agent that is adjacent to the client. The adjacent Bootp relay agent then sends the Bootreply to the client's station and IP addresses on the interface from which the Bootrequest was received.

Note that the Bootp relay agent functionality need not be turned on in all the routers between the client and server. The router that is adjacent to the client must be a relay agent. If that adjacent relay agent has a configured set of Bootrequest destinations that are either specific server addresses or subnet addresses, then it is the only router that needs the relay agent functionality. See figure 15. If the adjacent relay agent does not have either server or subnet addresses configured, then the Bootrequest is sent out through all of the adjacent router's interfaces, and is addressed to the All Hosts Broadcast for each interface. See figure 16. To allow these Bootrequests to be forwarded to the server, each router in the path between the client and server must be a Bootp relay agent.

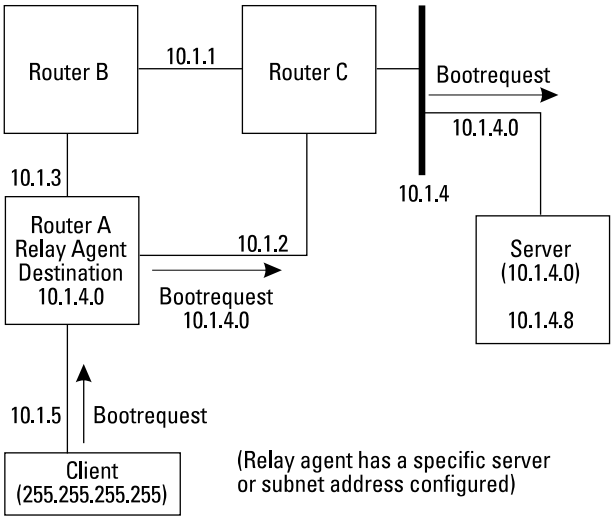

Figure 15. Bootrequest Relay with Specific Server or Subnet

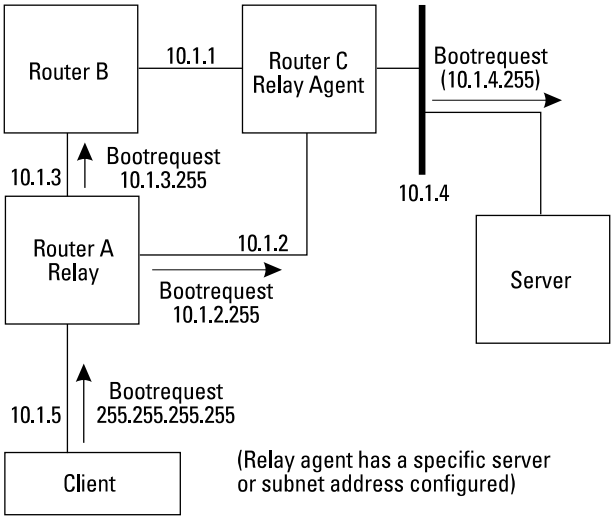

Figure 16. Bootrequest Relay with No Specific Server or Subnet

The Bootp relay agent can be configured with a list of destination addresses for Bootrequest packets. An address on this list can be any one of three types, as follows. The relay agent generates one or more Bootrequest packets for each address on the list, depending on the address type.

Specific Host Address is the address of a Bootp server somewhere on the network. The relay agent sends the Bootrequest packet to the specific host address. In this case, the Bootrequest is routed through the network to the Bootp server just like any other IP packet. No other relay agents are needed other than the router adjacent to the client.

Subnet Address is an IP address with a valid value for the network number and subnet number, but either all zeroes or all ones for the host part. For example,, the class A network number 15 with a subnet mask of 255.255.255.0 has the valid subnet address 15.1.1.0; 15.1.1.255.

The address 15.1.255.255 would not be a valid subnet address. This would be treated as a broadcast to all subnets in network 15. If a subnet address is configured as a Bootrequest destination,, the relay agent will set the destination IP address of the Bootrequest to the subnet address. The IP layer in the router will send the packet out whichever interface provides the shortest path to that subnet. The Bootrequest will be routed through the network to the subnet just like any other IP packet. No other relay agents are needed other than the router who is adjacent to the client.

All-Networks or All-Subnets Broadcast Address
is 255.255.255.255, the default used when there are no addresses configured in the Bootrequest destination list. An all-subnets broadcast address has the network number followed by all zeroes or all ones. For example, a network 15 all-subnets broadcast is 15.255.255.255. The class B network number 128.1 has an all-subnets broadcast of 128.1.255.255.

There are two cases to consider: either the relay agent has an interface in the specified network, or the relay agent does not have an interface in the specified network. If the router has one or more interfaces in the specified network, it transmits a Bootrequest packet on each adjacent subnet with the specified network number. The destination IP address of each packet is the all-hosts broadcast address for the subnet on which the packet is sent (for example, 15.1.1.255). Since the packets are addressed only to the adjacent subnet, all routers in the path between the client and server must have Bootp relay agent functionality. For the second case, routers that are not relay agents drop the Bootrequest packets. The other relay agents relay the Bootrequest packets in accordance with their configured destination address lists.

If the router does not have an adjacent subnet in the specified network, it will forward the Bootrequest along the path to the destination network as learned by the routing protocol. Thus,, no relay agents are needed between the relay agent that initiated the Bootrequest and the first router with an interface in the destination network. The first router that has an interface in the destination network must be a relay agent. It forwards the Bootrequest in accordance with its configured destination address list.

Virtual IP Host on Non-IP Networks

A router is often used in non-IP environments, for example, as a Novell IPX router or as a bridge. Sometimes in these environments it is desirable to have access to the router device management capabilities listed above, but without the burden of planning and configuring an IP routing network. The HP routers provide this capability by supporting an “IP host-only bridging” mode (as an alternative to the “router and host” mode for IP routing), in which the router is configured as a virtual IP host to support Ping, Telnet, TFTP, time protocol, and SNMP management. Routing protocols are disabled in this mode, and IP packets are bridged, not routed, throughout the network.

To use the device management capabilities of the routers, as well as those of the network management station, IP can be configured in IP host-only bridging mode. In this mode, the routers will not send routing messages of any type. IP addresses must be assigned only to router interfaces on which you anticipate having device management (IP) traffic information. The example network in figure 17 is configured to have IP traffic on all links except the LANs in Chicago and Washington D.C. Note that the IP addresses are the same for all IP interfaces on any single router. The subnet mask is the same for all IP router interfaces and devices throughout the network. The same IP address resolution protocol must be used on all of the IP router interfaces—ARP, HP Probe, or both.

Note that the bridging service must always be enabled to support this IP host-only bridging mode. All IP network traffic will be bridged. In the bridge network, there is no need to divide the IP network into subnetworks since the entire network is now a single IP subnetwork, A.

Consider the Novell IPX internetwork in figure 17. The three sites are each connected to one another through a router over point-to-point WAN links. Novell IPX traffic can be present on all links that are shown. On the LAN at the Los Angeles site is a network management station running the SNMP-based HP OpenView Interconnect Manager application.

Suppose that there are network-layer-transparent devices, such as repeaters or bridges, that support SNMP management on the LAN in Chicago. To access these devices from the router or the network management station, bridging and IP services would have to be enabled on router 2's interface to that LAN. An IP address, 1.0.0.2, and subnet mask, 255.0.0.0, would be assigned to that interface. The devices would be assigned unique IP addresses on network 1.0.0.0 and would also be assigned subnet mask 255.0.0.0.

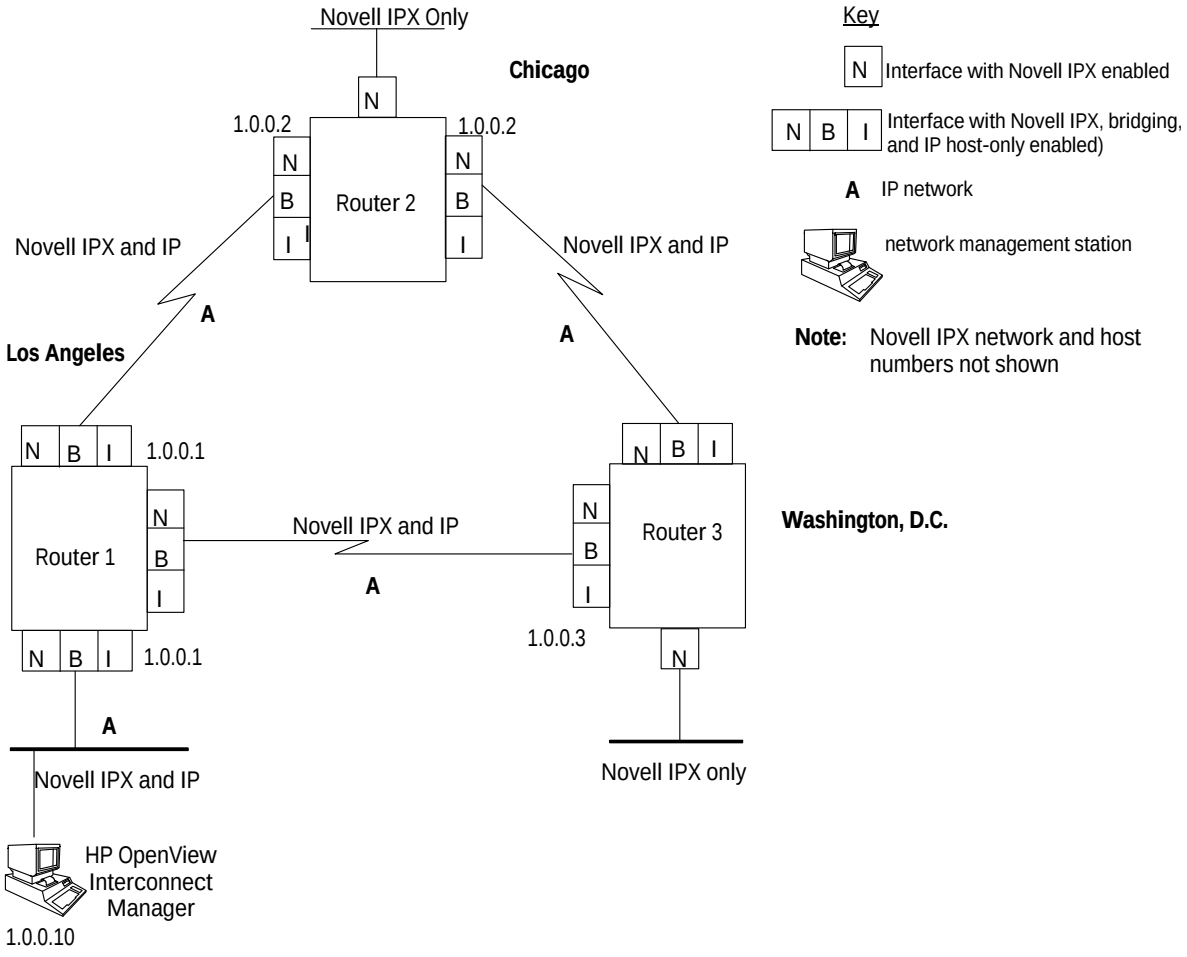

Figure 17. Virtual IP Host Configuration on a Novell IPX Internetwork

In host-only mode the HP router acts as an IP end node if bridging is not configured, or as a bridge if bridging is configured to a network management station. When using SNMP/IP-based HP Openview network management products, the router is not autodiscovered and displayed on maps as a router. This is because it does not meet the criterial for a router, even though its ID is for a router. Also, because there is no IP routing information being kept, the router cannot act as a default gateway for HP management products to provide routing information for autodiscovery.

In an IPX network, the router is described by an SNMP/IPX Openview DOS platform as a router. However, the router only supports SNMP/IP for accessing the MIB.

Source Routing and Token Ring Support

There are two ways traffic may communicate in an internetwork mixed with token ring/IEEE 802.5 and Ethernet/IEEE 802.3 networks. See the "Briding Service" note earlier in this manual for a detailed discussion of source routing and support for token ring LANs.

- At the bridge layer (also called the MAC layer), source-routing bridging is used to bridge token-ring-to-token-ring traffic, including traffic on the other side of source-routing bridges. Source-route *translational* bridging is used to bridge Ethernet/IEEE 802.3 and token ring/IEEE 802.5 LANs.
- At the routing layer (also called the network layer), the IP routing service provides source-route end-node support. Source-route end-node support allows IP nodes on Ethernet/IEEE 802.3 and token ring/IEEE 802.5 LANs to communicate with each other. You should enable source routing when you are configuring the IP routing services for each of your router's network interfaces.

IP Router Operations

Use the Network Control Language Interpreter (NCL) to display IP events, access IP statistics, view ARP and IP tables, and test for reachability of IP nodes. IP and SNMP must be enabled either in router mode or host-only bridging mode in order to access bridging address tables and routing tables for any of the routing services.

For detailed information on the router operations available in the Network Control Language Interpreter, see the reference manual for the HP routers.

Specifications

The specifications for the IP protocols are documented in a numbered series of technical reports called “Requests for Comments” or RFCs. The Internet Activities Board is the official committee for IP-related standards issued as RFCs. RFCs are distributed by the Network Information Center (NIC) at Government Systems, Inc., and can be obtained by postal mail or directly across the Internet using a file transfer program.

The following RFCs describe specific protocols and functions of IP routing. The IP routing service on HP routers is compatible with them:

General	IP	RFC 791
	TCP	RFC 793
	UDP	RFC 768
	Subnetting	RFC 950
	Gateways	RFC 1009
	IEEE 802 networks	RFC 1042
	MTU discovery	RFC 1063
Routing Protocols	RIP	RFC 1058
	OSPF	RFC 1247
	EGP	RFC 904
	Static routes	—
Address Resolution	ARP	RFC 826
	HP Probe	HP proprietary
	IP over X.25	RFC 877
Filtering	Source address	—
	Destination address	—
	TCP/UDP port number	RFC 1010
	Import route	—
	Export route	—
Device Management	Ping (ICMP)	RFC 792
	Telnet	RFC 854
	TFTP	RFC 783
	Bootp & DHCP	RFC 951 RFC 1084 RFC 1533 RFC 1534 RFC 1541 RFC 1542
	Time protocol	RFC 868
	SNMP agent	RFC 1155 RFC 1156 RFC 1157

Novell IPX Routing Service

Novell NetWare LANs are generally PC and/or workstation environments. NetWare supports a wide variety of LAN topologies and media. The HP routers support the Novell Internetwork Packet Exchange (IPX) routing service and the Sequenced Packet Exchange (SPX) protocol. IPX is the network-layer communication protocol used by Novell NetWare. Using HP routers you can interconnect NetWare LANs over a variety of WAN links to form an IPX internetwork. NetWare clients can then access servers on both local and remote networks.

Novell routers use the IPXWAN negotiation protocol over their WAN links. IPXWAN was described in RFC 1362 by M.Allen,, “Novell IPX over Various WAN Media (IPXWAN)”, Novell Inc., September, 1992. A successor, IPXWAN Version 2 (IW2), is now used. IW2 specifies initial connection setup methods for running IPX traffic over WAN media such as PPP, X.25 switched virtual circuits, X.25 permanent virtual circuits, and frame relay. HP routers will use the IPXWAN (and IW2) protocol when it is configured for a WAN link, to interoperate with a Novell router.

NetWare Services

Novell NetWare provides services that vendors and end users can use to develop distributed applications. The basic services in a Novell network are a collection of functions provided by a file server. The protocols used to obtain these services are called the NetWare Core Protocols (NCP). They are implemented at the ISO application layer on Novell end nodes (clients and servers), but not on routers. ISO transport-layer functions are handled by the Packet Exchange Protocol (PEP), which is often documented as a part of NCP in Novell literature. A list of services commonly provided by Novell NetWare follows.

- *Bindery Services:* manage name resolution, accounting, and security.
- *Diagnostic Services:* retrieve software setup and status.
- *Directory Services:* access a distributed file system directory.
- *File Services:* access remote files.
- *Message Services:* deliver datagrams to IPX nodes.
- *Message Handling Service* (MHS): handle electronic mail file and exchange formats and procedures
- *Queue Management Services* (QMS): handle control queues of incoming requests, including processing and printer queues.
- *Synchronization Services:* coordinate tasks among multiple users, including file and record locking and semaphores.
- *Transaction Tracking System* (TTS): ensure the integrity of file operations in the case of system failures.

General Addressing Considerations

An IPX packet header contains a source network address field and a destination network address field. The network address is composed of three fields, as shown in figure 1. The network number uniquely identifies a network in the internetwork. A host number uniquely identifies a single host. The socket number uniquely identifies a socket within the operating system of a host. A socket is the address of a higher-level protocol—Routing Information Protocol (RIP) for example—that is using the services of IPX. The source and destination network addresses uniquely define the communicating sockets in the entire internetwork.

When IPXWAN is used, the IW2 header is contained within the data portion shown in figure 1.

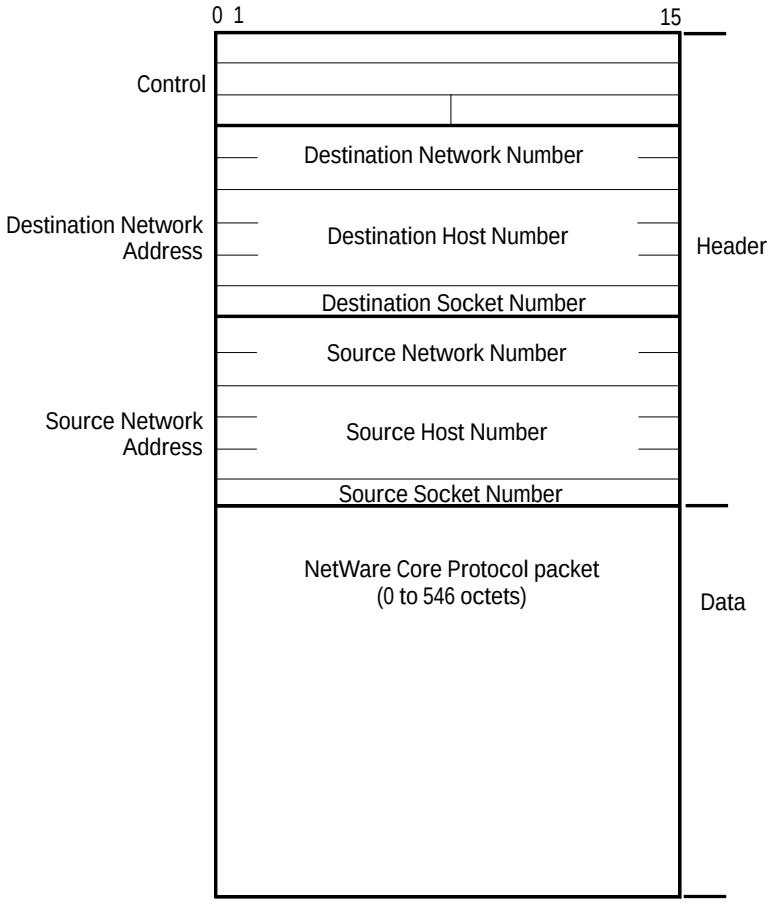

Figure 1. IPX Packet

Planning for an IPX internetwork involves assigning a unique host address to each end node and router, and assigning a unique network number to each LAN and WAN link. Socket numbers are not planned, but are dynamically built into the IPX packet header by the IPX protocol before the packet is sent to the data-link layer.

IPX host addresses are 48-bit numbers (12 hexadecimal digits). For many IPX nodes, as well as HP routers, the station address (also called MAC or physical address) of the device is used as the host address. This ensures uniqueness and is one less parameter to configure. For the HP routers you do not specify an IPX host address; the station address of WAN port 1 will be used (except for operating system versions earlier than 5.70).

IPX network numbers are 32-bit numbers (8 hexadecimal digits). On HP routers, a unique network number must be assigned to each network interface where IPX routing service will be enabled. Devices connected to the same network as a router interface must use the same network number that is configured on that router interface. Network numbers in the range 00000001h to FFFFFFFEh are valid assignments. The network numbers 00000000h (all zeros) and FFFFFFFFh (all ones) are reserved to mean unknown network and all networks, respectively.

Figure 2 shows an IPX internetwork that connects subnetworks in five cities using HP routers with the IPX routing service enabled. When configuring the IPX routing service, you will assign IPX addresses like those shown to the router.

The IPX address consists of a network number and a host number. The network number can be an 8-digit hexadecimal number. Subnetworks that are connected by bridges are assigned the same network number. Subnetworks that are connected by routers must be assigned different network numbers. That is, an 8-digit hexadecimal number must be assigned to each router LAN (Ethernet/IEEE 802.3 and/or token ring) and each WAN port that will be routing IPX packets. For example, in figure 2, the HP router in Los Angeles connects an Ethernet and a token ring LAN. Each LAN is assigned a separate network number, 00000004 and 00000005, respectively. The WAN link that connects Los Angeles to San Francisco is also assigned its own network number, 00000003.

Each end node is assigned to an IPX network and has a unique host number. The IPX host number is the station address assigned by the network equipment manufacturer. You don't need to assign or configure the host number; each node on the network has a unique host number.

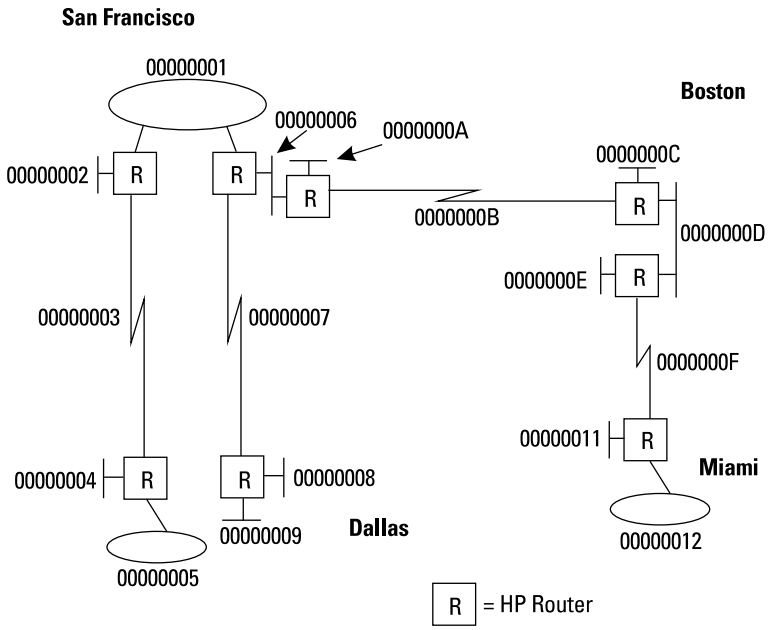

Figure 2. IPX Internetwork

Novell Data-Link Layer Encapsulation

In addition to Ethernet and IEEE 802.2 encapsulation for data-link-layer frames, IPX routing service supports the proprietary Novell encapsulation. Novell encapsulation, as shown below, prefixes an eight-octet preamble, six octets of destination-address information, six octets of source-address information, and two octets of packet-length information to the IPX packet. It appends a four-octet frame check sequence to the packet. Choose Novell encapsulation when using HP routers in Novell NetWare internetworks.

Note

This encapsulation choice will be valid only for IPX routing services and will have no effect on the encapsulation used by other enabled bridging or routing services.

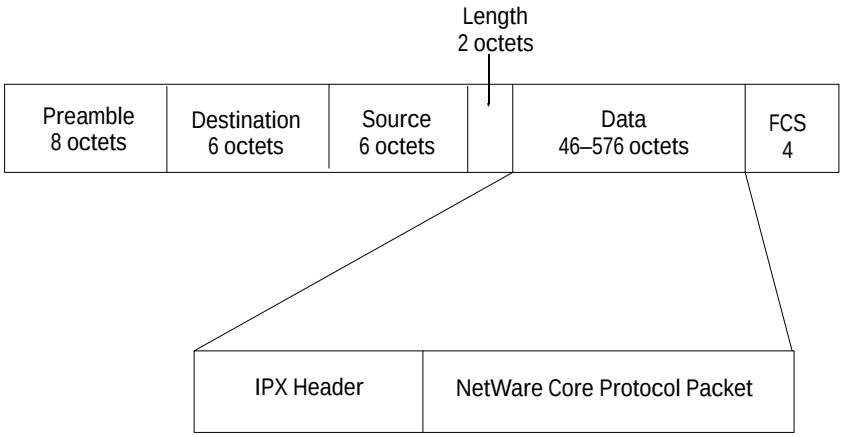

Figure 3. NovelI Encapsulation

The maximum length of an IPX packet passed through the router is 576 bytes. Note that on an Ethernet you may see larger packets. If these packets are destined for remote LANs through the router, the data will be truncated before the router forwards it.

IPX Routing Table

The IPX routing service forwards packets based on:

- the network number portion of the destination network address, and
- the information in the IPX routing table.

The IPX routing table contains an entry for each network number learned by the router. For each destination network number entry, the table also contains the following information:

- *Next Hop:* If the destination network is directly connected to this router, then next hop is the host number of this router (station address of WAN port 1). If the destination network is a remote network, the next hop is the host number of the next router along the path to the destination.
- *Hop:* the number of hops to the destination network from this router. “n” hops indicates that “n” routers separate this router and the destination network.

- *Route Type* specifies whether the destination network is remote or directly connected. Also, if the router is notified that a learned route is no longer available, then the route type is marked invalid. The route will remain in the table until the router is notified of another valid route to the destination or the router is rebooted.
- *Route Learned* indicates whether the route was statically learned (configured) or learned dynamically by the Routing Information Protocol (RIP).
- *Age* indicates the number of seconds since the route was learned.
- *Interface* is a number, assigned by the router, indicating the interface through which the next hop is reached.

For more information on the IPX routing table, please refer to the operator's guide for an HP router.

Note

While the IPX protocol uses the information in the routing table to determine where to deliver each packet, the IPX protocol does not gather the information to create and maintain the routing table. That function is performed by RIP. Alternatively, routes can be statically configured.

Routing Information Protocol

The Routing Information Protocol (RIP) is the dynamic routing protocol used in Novell IPX internetworks for gathering and maintaining routing information. Novell RIP supports both request and response operations. RIP requests are used by hosts to determine the network number of the network that they are connected to or to determine the route to a specific network. A router or another host can generate a response to a RIP request.

An IPX router periodically generates a RIP update message as an unrequested response operation. A RIP update message contains all of the internetwork topology information that the sending router has in its routing database (table). When a router detects that a network is unavailable or that its route to that network has changed, that information will be reflected in its next RIP update message. HP routers transmit RIP update messages on all IPX interfaces every 30 seconds. This transmission period is not configurable on HP routers. However, you can disable the generation of RIP update messages on each interface.

RIP requests and responses are sent as data in IPX packets. The Novell NetWare version of RIP uses IPX socket 453h.

RIP imposes a network diameter limit of 15 hops or less on IPX internetworks. To enable RIP to account for different link speeds, HP routers allow you to assign a higher cost to interfaces for lower-speed networks.

IPX Static Routes

For IPX routing service, you can also configure static routes (not to be confused with NetBIOS static routes described later). Static routes are used to restrict paths that IPX packets can follow to those that you define. Static routes will be displayed in the IPX routing table and cannot be overwritten. When you configure a static route, you must configure the following parameters for each destination network:

- target (destination) network
- next-hop host number
- next-hop network number
- RIP table cost (hop-count value)

For more information on configuring IPX static routes, refer to the IPX configuration parameters in the reference manual.

Service Advertising Protocol

The Service Advertising Protocol (SAP) is used by NetWare servers to inform NetWare clients of their presence. SAP is an ISO application-layer protocol that runs directly over IPX, an ISO network-layer protocol. The Service Advertising Protocol allows a file/print server or program to register its name on the network. The programs may be application programs written by third-party developers or users. A SAP request will ask for the translation of that name to a socket on a particular node on the network. The NetWare Core Protocols (NCP), used to obtain the core service offered by a NetWare file server, use SAP to find file servers or print servers in the internetwork.

Servers broadcast service advertising packets every 60 seconds. These packets identify a server by name, server type, and network address (network, host, and socket identifiers). All routers receive these packets. The HP routers maintain a SAP bindery (table) of server information (name, type, address, hop count, interface to server, and timer value) based on the Service advertising packets they receive. HP routers update the age timer for existing entries. If new server information is discovered, the router creates

or changes a bindery entry and broadcasts service advertising packets to other networks. On HP routers, the SAP bindery (table) can be displayed using NCL's Rgetis command. For more information on the IPX SAP table, refer to the operator's guide for an HP router.

HP routers have a service advertising socket (SAS)—452h—used to respond to client requests for information on network servers. HP routers also broadcast unsolicited SAP packets, called GSRs (general service responses). GSRs contain the entire bindery if no SAP filters are configured. GSRs are broadcast every 60 seconds on LANs. On WAN circuits, you can configure the GSR broadcast period in the range of 1 to 99 minutes, or you can disable GSR broadcasts.

NetWare clients use the IPX broadcast facility to obtain information on network servers. Two types of queries are supported. General-service queries solicit information on all network servers. Nearest-service queries solicit information on the closest service of a specific type.

SAP Filters

HP routers support SAP filters, which are used to logically partition IPX internetworks by controlling the advertisement of servers. A NetWare client cannot access a server that is not advertised to it. All servers in the bindery are advertised if no filters are configured.

SAP filters are configured for each interface. Two levels of SAP filters are supported. Network-level filters permit or deny advertising of servers that match a pattern of network number and server type on a given interface. Server-level filters allow you to permit or deny advertising servers that match a pattern of server name and server type on a given interface. Server names can be up to 48 case-sensitive characters.

For both filter levels, there are patterns to indicate all networks or all server types. Up to 50 filters of each level can be configured. When both levels of filters are configured on a single interface, server-level filters take precedence over network-level filters. Also, for multiple filters of a given level, the ones that are configured first have lower precedence than those that are configured later.

For more information on configuring SAP filters, refer to the IPX configuration parameters in the reference manual.

NetBIOS Protocol Support

The Network Basic Input/Output System (NetBIOS) is a widely implemented session-layer protocol developed by Sytek, Inc., for IBM PC networks. Many vendors have written programs that are compatible with NetBIOS. NetBIOS broadcasts are used by client programs to establish the connections with servers.

NetBIOS was originally designed to directly use services provided by data-link-layer protocols. To allow client programs to run across an IPX internetwork, Novell NetWare provides a NetBIOS emulator. Novell NetBIOS runs over the Packet Exchange Protocol (PEP), at the transport layer, and over IPX, at the network layer.

IPX routing service on HP routers can be configured to either forward or restrict NetBIOS broadcast packets. NetBIOS broadcast packets received on one of the HP router's IPX interfaces will be broadcast on all other IPX interfaces using the IPX "all nets" broadcast facility. The acceptance and/or rebroadcasting of NetBIOS broadcast packets can be configured for each interface. Therefore, the router has the capability to restrict server or client access to remote connections.

There are two parameters used to configure Novell NetBIOS broadcast handling on an IPX router interface. Set the Accept NetBIOS Bcasts from Net parameter to enable or disable client access on a given IPX network to NetBIOS servers on other networks. When enabled, client-generated NetBIOS broadcast will be sent to other networks. Set the Deliver NetBIOS Bcasts to Net parameter to enable or disable access to servers on a given IPX network. When enabled, NetBIOS broadcasts from other networks will be sent to this network.

Novell NetBIOS can be used in IPX internetworks with up to a maximum network diameter of eight hops. However, this restriction, as well as problems related to excessive NetBIOS broadcast traffic, can be overcome by using NetBIOS broadcast static routes.

For more information on configuring NetBIOS broadcasts, refer to the IPX configuration parameters in the reference manual.

NetBIOS Broadcast Static Routes

HP routers provide a non-Novell-standard static-routing mechanism that converts IPX “all nets” NetBIOS broadcast packets to directed broadcast packets. A directed broadcast is an IPX network-level broadcast to a single network. NetBIOS broadcast static routing is typically used to enable a NetBIOS client of one network to establish a session with a remote NetBIOS server. The ability to control session establishment facilitates internetwork security and management. Additionally, use of static routes can significantly reduce the amount of NetBIOS broadcast packets in the IPX internet.

To configure a NetBIOS static route, you must specify a NetBIOS resource name and the destination network where it resides. If the NetBIOS resource name in a broadcast packet matches a static-route-table entry, then the NetBIOS packet is routed to the destination network. If no match is found, then the packet is treated as specified by the Accept NetBIOS Bcasts from Net and Deliver NetBIOS Bcasts to Net parameters. Up to 50 NetBIOS broadcast static routes can be configured.

Note

Since the NetBIOS static route is not a standard Novell feature, it may not work when non-HP or non-Wellfleet routers are used in the IPX internet.

Refer to the IPX configuration parameters in the reference manual.

Novell IPX Routing Service
NetBIOS Protocol Support

AppleTalk Phase 2 Routing Service

Product Note

The HP routers support AppleTalk Phase 2 routing over Ethernet or token ring links and synchronous WAN links. AppleTalk Phase 1 routing is not supported; however, it can be relayed using the bridging service. More detailed information about AppleTalk Phase 2 can be found in the configuration guide for an HP router.

An AppleTalk internet is a collection of one or more AppleTalk networks connected by AppleTalk routers. An AppleTalk network is a LAN that contains end nodes and connects to at least one AppleTalk router. The HP routers support CSMA/CD LANs (EtherTalk) and token ring LANs (TokenTalk).

There are three ways that routers are used to build an internetwork.

- *Local router:* A router used to connect LANs only is called a local router. The local router is the only router attached to the LAN.

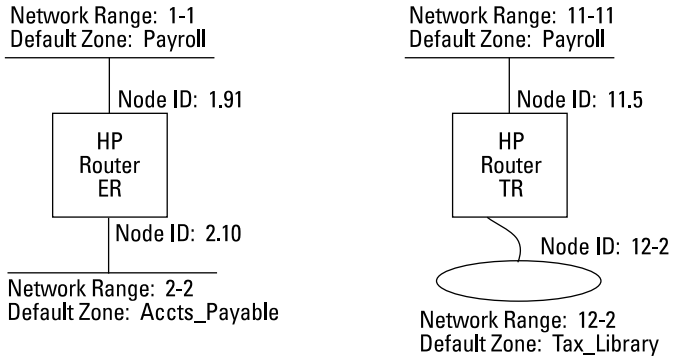

Figure 1. Local Router Configuration

- *Half router:* Two routers, each connected to one or more AppleTalk LANs, and connected to each other through long-distance communication links, are called half routers.

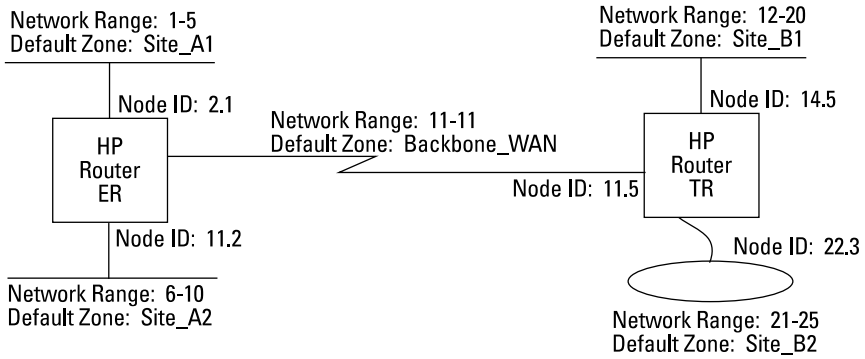

Figure 2. Half Router Configuration

- *Backbone router:* Routers, each connected to one or more AppleTalk LANs, connected together through either an Ethernet backbone or an X.25 packet-switching backbone network, are called backbone routers. Backbone networks do not have any AppleTalk end nodes attached.

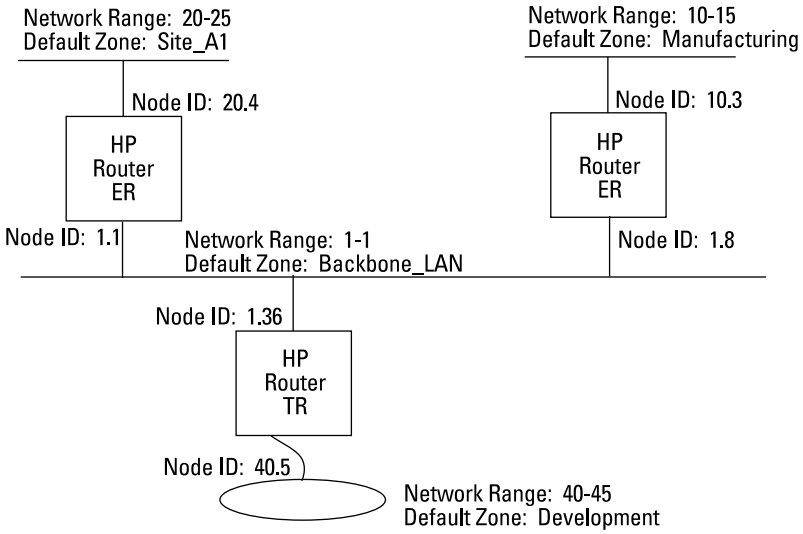

Figure 3. Backbone Router Configuration

A sample network map for an AppleTalk internetwork is shown below in figure 4.

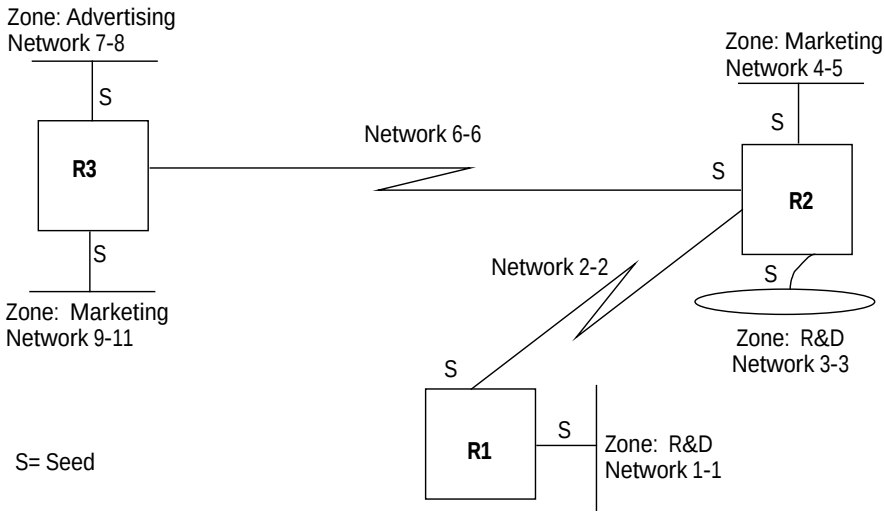

Figure 4. AppleTalk Network Map Example

Assigning AppleTalk Addresses

In an AppleTalk internet, nodes (end systems or routers) and networks are assigned addresses. Nodes and routers are each assigned node addresses that are unique throughout the internet.

The AppleTalk node address consists of:

- a 16-bit network number (1–65279)
- an 8-bit node identifier (1–253)

Each end node is assigned to an AppleTalk network and has a unique node identifier. Each AppleTalk network is assigned a range of network numbers. There can be up to 253 nodes for each number in the network number range. Each network number range must be unique and must not overlap.

This network number range is configured by the user on at least one of the routers residing on the network. That router is called the seed router for that network. The AppleTalk nodes dynamically determine the network number portion of their node address. If the node is being restarted, and there is a network number stored in non-volatile memory, that number will be used. If not, a network number will be chosen from the startup range configured by the user on the seed router, which is discussed in more detail below.

The node identifier portion of the node address is assigned by the AppleTalk software automatically.

Zones

AppleTalk nodes and networks are assigned zones in which they will reside. Zones are logical grouping of nodes that share the same network resources. Zones may encompass more than one network. Nodes in a zone need not be physically contiguous and need not have the same network number. A unique zone name, an alphanumeric string of up to 32 characters, is assigned to each zone.

If you assign easy-to-remember zone names, network devices in a large internet can be looked up easily. Name searching can be done within one or more user-specified zones. Each node in an AppleTalk internet can select one of several possible zone names from the list of allowable names for its network. The seed router on each AppleTalk LAN network is configured with a list of zone names that can be used by nodes on that LAN. If a node on a network is assigned a zone name that is not on the configured list, the seed router will assign the default zone name to that node. The default zone name is the AppleTalk zone name that will be assigned to any node in the attached network that was not assigned its own zone name.

Seed Routers

A router identified as a seed router for a network has been configured with the network number range and default zone name for all the nodes that reside on the attached network. (Both the network number range and the default zone name can be configured using either Quick Configuration or the Configuration Editor.)

At least one seed router, configured with the network number range and default zone name, must exist on each AppleTalk network or WAN link. More than one seed router can be assigned to a network for redundancy. However, all seed routers for a particular network must be configured with the same network number range and zone name information.

The seed router dynamically sends the network number to the other connected routers using RTMP routing protocol messages. If you do not configure the network number on a router, that router can learn the number from the seed router.

The seed router can optionally be configured with a list of zone names that will be valid on the network. Any nodes on the network with a node name not on this list will be assigned the default zone name by the seed router. (Note that the zone name list cannot be configured using Quick Configuration; it must be configured using the Configuration Editor.)

Since a router usually connects to more than one AppleTalk network, a single router may be the seed router for the network on one port but not for the network on another port. A router can also be configured as the seed router for different networks on different ports. Note that the network number range configured on each of the individual ports must be different and may not overlap.

AppleTalk Protocols

The AppleTalk protocols provide network-access standards for layers one through five (in the Open Systems Interconnection (OSI) reference model). Protocols for layers two through five are shown in figure 5.

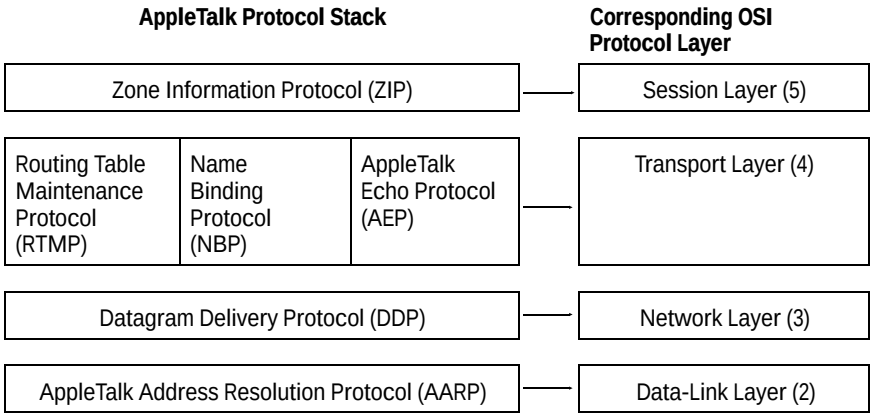

Figure 5. AppleTalk Protocols

AppleTalk Address Resolution Protocol (AARP)

- Translates node addresses into equivalent station addresses.
- Maintains the Address Mapping Table (AMT), which contains a listing of equivalent node and station addresses.
- Ensures the integrity of node addresses. When a node address is assigned, AARP confirms the uniqueness of that address.

When a node identifier is dynamically determined, it must be automatically validated as a unique node identifier on a given AppleTalk LAN. AARP generates a random node identifier, or, if the node is being restarted and a node identifier already exists in non-volatile memory, that node identifier will be tested. To ensure that the node identifier is unique, AARP compares it to entries in the Address Mapping Table. If no conflicts are found, then AARP transmits a series of AARP Probe packets to the tentative address. If there is no response, then AARP validates and uses the node identifier. AARP Probe can be disabled if explicit addresses are going to be used. However, it is recommended that dynamic addressing be used and that AARP Probe remain at the default, enabled.

Datagram Delivery Protocol (DDP)

- Provides a “best-effort” socket-to-socket delivery of datagrams over an AppleTalk internet.
- Acquires the AppleTalk network number.

Routing Table Maintenance Protocol (RTMP)

- Builds and maintains the AppleTalk routing table. Each table entry includes a destination network range, the AppleTalk node address (network number and node identifier) through which the destination is reached, the number of router hops to the destination, and the route status.

Zone Information Protocol (ZIP)

- Determines which networks belong to which zones by maintaining the Zone Information Table (ZIT) on each router. The ZIT provides an internet-wide mapping of network number ranges to zone names.

  Note that the ZIT is not displayed by the HP routers. However, the Local Zone Table, a table of zone names configured on a particular router, can be displayed using NCL's Rgetat command.

Name Binding Protocol (NBP) on a Router

- Maps network entity names with internet addresses. It allows you to specify descriptive or symbolic names, while other processes may refer to the same entity numerically (for instance, by node addresses).
- Facilitates the internet-wide device-name-lookup process.
- Uses ZIP to determine which networks contain nodes that belong to a zone.
- Provides a list of available services within a zone to clients.

AppleTalk Echo Protocol (AEP)

- Tests whether a destination node is reachable. AEP enables a node to send a request packet to another internet node and to receive a response packet back. The command used is Atping.

DECnet Routing Service

The HP routers support DECnet Phase IV routing services. Unlike IP, IPX, AppleTalk, and XNS routing services, source-routing support is not provided for the DECnet routing service. The HP routers implement the DECnet routing protocol (DRP), the network-layer protocol of the Digital Network Architecture (DNA). DEC systems support several other network architectures, including:

- *Local Area Transport Architecture* (LAT): proprietary DEC architecture for terminal servers on Ethernet networks. LAT protocol is a direct user of Ethernet service.
- *Maintenance Operations Protocol* (MOP): implemented on diskless nodes, such as the MicroVAX 2000, for downloading their operating systems. MOP is a direct user of the data-link layer services.
- *System Communications Architecture* (SCA): also known as "VAX clusters", uses System Communication Services protocol at the network layer. This architecture allows multiple computers to share disk drives over Ethernet at 10 Mbit/s or over the Computer Interconnect (CI) bus at 70 Mbit/s.
- *DECnet/OSI Architecture*: DEC's implementation of the Open Systems Interconnect (OSI) protocols for DECnet Phase V. It uses ISO (International Standards Organization) network-layer protocols.

It is not uncommon to see DECnet (DNA), VAX clusters, and LAT all providing services on a single Ethernet. Note that only DNA-compatible products—HP routers for example—implement DECnet routing protocol. The DECnet routing service on HP routers can only be used for communication between networks if the communicating devices are DNA-compatible. Otherwise, the bridging service must be used. HP routers will support OSI routing in the near future.

DECnet Services

The following are among the services available on a DECnet internetwork through HP routers:

- *Virtual Terminal Service*: allows a remote terminal to access a host system. (Also known as CTERM, control terminal module).
- *Data Access Protocol* (DAP): is a presentation-layer protocol suite that provides functions for exchanging data between two nodes on a network.
- *Record Management Service*: is a common I/O interface for VMS used for accessing data from a remote node. It uses DAP services.
- *Network File Transfer* (NFT): is an interactive utility providing access to remote data. It uses DAP services.
- *Distributed Naming Service* (DNS): is a directory service that separates the logical name of an entity from its physical location.
- *Distributed File System* (DFS): allows a remote file to appear to the user as though it were local.
- *Distributed Queuing Service*: provides access to remote printers.
- *Videotex* is an interactive system that allows subscribers to down-load several pages of information to a local terminal or node.

General Addressing Considerations

DECnet Phase IV uses a hierarchical addressing scheme. DECnet internetworks can be logically divided into distinct, non-overlapping areas. An area can be a single LAN or several LANs and WANs interconnected through routers. A DECnet internetwork can be divided into as many as 63 areas. An area can have as many as 1023 nodes (routers and end nodes). End nodes within an area reside only on the LANs. Routers connect both LANs and WANs and may connect to routers in other areas. Note that for performance reasons, all nodes on a LAN must be
in the same area.

Dividing a DECnet internetwork into areas provides two advantages. First, it takes advantage of the entire range of network addresses available to you. Second, it improves the efficiency of your internetwork by reducing the volume of control messages. Each router maintains a database of all nodes that are in its area. This information is periodically exchanged with other routers in the area. The only information that is maintained about other areas is the route to take to reach those areas. Dividing the internetwork into areas reduces the amount of information that routers have to maintain and exchange with one another.

Planning for a DECnet internetwork involves defining areas and assigning unique addresses to each router and end node in the internetwork. When configuring DECnet nodes, individual addresses must be configured on each router and end node. Note that a single address is assigned to the router itself, not to each port or interface on the router.

DECnet addresses consist of:

- a 6-bit area number (1–63)
- a 10-bit node number (1–1023)

The area number and the node number are separated by a period. For example, the DECnet address 1.10 specifies node 10 in area 1.

A DECnet internetwork is shown is figure 1. It consists of two areas connected through an X.25 network.

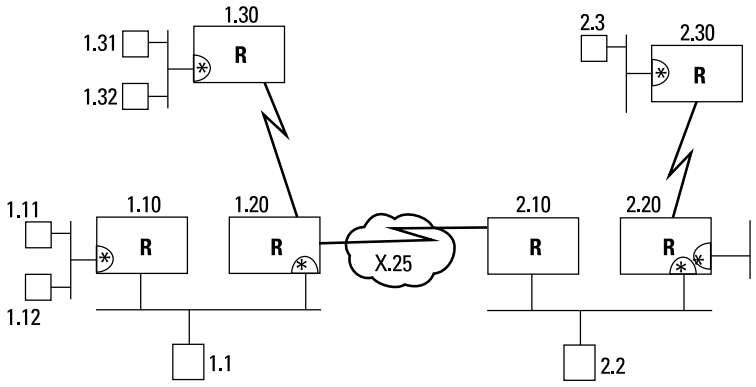

Figure 1. DECnet Internetwork

Station Address Resolution

On each DECnet node, network-layer addresses are resolved to station (MAC) addresses when the node initializes. The station address on these nodes will be modified based on the area number and node number that you have configured on the nodes. This also happens for the network interfaces (circuit groups) on the router. The station address will always have the number AA0004h, the DEC vendor code, as its most-significant 24 bits. The uniqueness of the network-layer addresses preserves the uniqueness of the station address.

The station address resolution for DECnet node 1.10 is shown in figure 2. The station address for node 1.10 will be AA0004000A04h.

	Decimal	Hexadecimal	Binary
DEC Vendor Code (24 bits):	—	AA0004	—
Area Number (6 bits):	1	—	000001
Node Number (10 bits):	10	—	0000001010

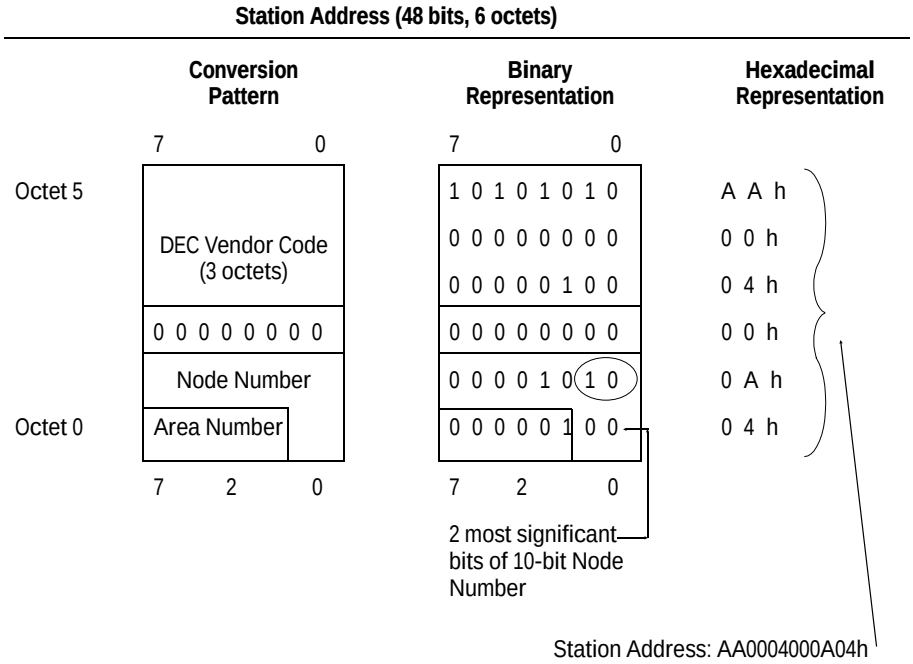

Figure 2. Station Address Resolution for Node 1.10

Note: When the DECnet routing service is enabled on HP routers, the station address will be modified, as described above, on all interfaces, even those where DECnet routing service is not enabled. This “DEC” station address will become the node ID for the router for IPX and XNS routing services. It takes precedence over an XNS node ID that you may have configured. An exception is on synchronous pass-through interfaces where the “Local LAN Address” that you configure will take precedence only on the pass-through interface.

DECnet Link Cost

DECnet network links are not assigned addresses. Instead, each link must have a cost value associated with it. Ideally, the cost value should reflect the speed of the link. The router uses this information to select the fastest (least-cost) path to a destination. Therefore, cost values have a significant effect on what routing decisions are made by the router.

Planning a DECnet internetwork involves determining the cost value for each link. These values should be documented on the network map and noted on the site survey worksheets. To configure link costs on HP routers, use the Configuration Editor. Quick Configuration will assign the default value if no other value has been previously configured. A table of recommended values as well as a the default cost value, “Suggested DECnet Circuit Costs”, is in the DECnet chapter in the reference manual.

Designated Router

DECnet end nodes always reside on LANs, and each maintains an "on-Ethernet cache" of other nodes on its LAN. When an end node sends packets destined for a node on a remote LAN (a node that is not in the cache), the packets are sent to the "designated router". The designated router routes packets on behalf of the end nodes. On each Ethernet, one router is elected as the designated router. If the end node has not learned the identity of the designated router, usually because no routers are present on that LAN, the end node will attempt to send the packet directly to the destination node.

When there are multiple routers on a LAN, one will be elected as the designated router. To control which router is elected, configure the Router Priority parameter on each network interface on each router. The router that is assigned the highest priority relative to the other routers on the LAN will be elected the designated router for that LAN. If there are routers with equally high priority, the router with the highest node number will be elected as the designated router. Note that a router that is connected to multiple LANs may be the designated router on some LANs but not on others. A designated router will be elected even if there are no end nodes currently on the Ethernet. In figure 1, designated router ports for each Ethernet LAN are highlighted with an asterisk (*).

Adjacency and Initialization

Before an end node can send packets through a DECnet router, the node must establish an adjacency with that router. Additionally, routers must establish adjacencies with one another before they can transfer packets between them. All nodes (end nodes and routers) on a LAN with a router establish adjacencies with that router. Two routers connected through a point-to-point WAN link establish an adjacency. Two or more routers connected to a common X.25 network (using X.25 point-to-point circuits) establish adjacencies with one another. Routers maintain a database of adjacent end nodes and routers. This is some of the information that is represented in the DECnet routing tables.

DECnet nodes establish adjacencies as part of their initialization sequence. End nodes multicast “Endnode Hello Messages” to establish adjacencies with routers on the same LAN. Routers multicast “Router Hello Messages” to establish adjacencies with one another and to elect the designated router. Additionally, the designated router multicasts “Router Hello Messages” to all end nodes to inform them that it is the designated router.

Hierarchical Routing

DECnet routing is hierarchical. Level 1 routing occurs when packets are routed within an area. Level 2 routing occurs when packets are routed between areas. An HP router always performs as both level 1 and level 2 routers, even if it is not adjacent to a router in another area.

A DECnet router maintains routing tables for both a level 1 and a level 2. The level 1 routing table contains information on routes to destination nodes within the area. The level 2 routing table contains information on routes to destination areas.

Figure 3 shows a two-area DECnet internetwork. Each area has two routers. When end node 1.10 sends a packet to end node 2.10, it is routed through all four routers in the internetwork. When router 1.1 receives the packet destined for node 2.10, the router must refer to its level 2 routing table to determine where to send the packet because the destination node is in a different area. It sends the packet to router 1.2. This router must also refer to its level 2 routing table to determine where to send the packet. It sends the packet to router 2.1. Since the destination node 2.10 is in the same area as router 2.1, the router refers to its level 1 routing table to determine where to send the packet. It sends the packet to router 2.2, which forwards it to end node 2.10.

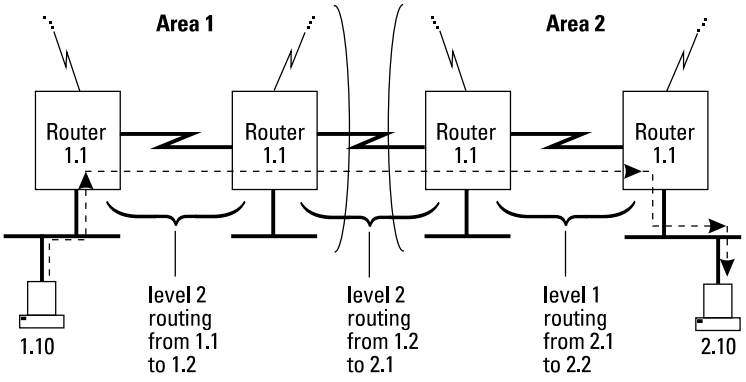

Figure 3. Hierarchical Routing of a Packet Sent from Node 1.10 to Node 2.10

DECnet Routing Metric

In both the level 1 and level 2 routing tables, the routing metric that is maintained is (path) cost. Each DECnet LAN and WAN has a cost value associated with it that relates to the speed of the link. Faster links have lower cost values associated with them. To plan the cost assignments for a DECnet internetwork and configure them on HP routers, see the “DECnet Link Cost” section above.

The path cost is the cumulative cost value of the links along the path to the destination node (in the level 1 routing table) or area (in the level 2 routing table). For a given destination node or area, DECnet routing tables reflect the least-cost path that the router has learned from periodically exchanging routing update messages.

Routing Table Maintenance

DECnet routers maintain their routing tables by multicasting routing messages to adjacent routers. A level 1 routing message is periodically multicast to adjacent routers in the same area. The level 1 routing message contains the sending router's current information on the routes to all nodes in the area. A level 2 routing message is periodically multicast to adjacent level 2 routers whether or not they are in the same area. The level 2 routing message contains the sending router's current information on the routes to all areas. When a router receives a routing update message, it may or may not update its routing table based on whether the router learns of a new destination node or area or a lower-cost path.

On HP routers, routing messages are multicast every 3 minutes by default. However, a period as short as every 15 seconds can be configured by modifying the Bcast. Routing Timer parameter using the Configuration Editor.

Connecting to Routers from DEC

HP routers use the HDLC data-link layer protocol when connected together over point-to-point links. DEC routers use the DDCMP data-link layer protocol over point-to-point links. These protocols are not compatible. Therefore, you cannot connect DEC routers to HP routers over a point-to-point serial link. You must use either an Ethernet or an X.25 link.

DECnet Router Operations

Use the Network Control Language Interpreter (NCL) to display DECnet events, access DECnet statistics, view DECnet routing tables, and display DECnet management information base (MIB) variables. For detailed information on NCL and the various DECnet routing service operations, refer to the user's guide and reference manual.

Specification

DECnet routing service on HP routers is compatible with the following specification:

- DECnet Digital Network Architecture Phase IV Routing Layer Functional Specification, December, 1983.

A Primer on HP Probe

HP Probe is a Hewlett-Packard proprietary protocol used on HP nodes. It is an unreliable-connectionless request reply protocol designed to provide the name-to-address mapping information between HP nodes using Network Services (NS), and on HP Data Communications and Terminal Controllers (DTCs). NS and some DTC services use TCP/IP as the transport, and therefore use IP and link-level station (MAC) addresses to address packets between nodes. However, NS users access nodes in terms of their names. On an NS network, a user supplies a node name to the local transport layer when he wants to set up a connection with it. (NS uses the format `node.domain.organization` for its node names.) If it does not already know the address for the destination node, then one or more Probe requests are generated. The destination of the request is a target node or a Probe name server (called a proxy server), which will respond with a Probe reply that contains the requested path or address information. A Probe name server is a machine that contains a mapping of names to addresses for the network and other nodes in the internet.

The information returned in a Probe reply packet is a path report. It contains the supported protocol stacks (e.g., TCP, IP, IEEE 802.2), services (e.g., DSCOPY), and address information of the target node. The path report is used by the requester to provide name-to-address mappings in order to create a connection or communicate with the target node.

HP Probe Protocol Definition

HP Probe supports both Ethernet and IEEE 802.3 encapsulation. HP Probe uses two multicast addresses: a primary multicast address 0x090009000001, and a secondary proxy multicast address 0x090009000002. Their use will be explained below. The HP Probe protocol uses 0x8005 as the Ethernet type and 0xFC for the IEEE 802 SAP. FC is HP's expansion XSAP indicator; the XSAP for Probe is 0x0503.

The following HP Probe packets are discussed below:

1. Probe Unsolicited Reply *
2. Probe Name Request
3. Probe Name Reply
4. Probe Virtual Address Request (VNA)
5. Probe Virtual Address Reply *
6. Probe Proxy Request
7. Probe Proxy Reply

* Supported by HP routers

Probe Unsolicited Reply

The Unsolicited Reply is transmitted over the primary multicast address when a node first comes on line and periodically afterwards. The Unsolicited Reply packet contains complete, updated information about a node. Included in the packet are the well-known NS services it supports, the protocols it supports, its name, and machine addresses. A node uses the Unsolicited Reply packet to announce its presence on the network or to update information about itself that was transmitted earlier and has changed. HP routers send an Unsolicited Reply upon booting, but do not listen for or receive Unsolicited Replies sent by other nodes.

Probe Name Request/Reply

When an NS LAN node (node A) wants to connect with another node (node B), node A first checks to see if it has a name-to-IP-address map in a database called a Nodal Registry. If it does not, it generates a name request packet containing the target node's name (node B), and sends it out over the primary multicast address. All nodes receive the packet, but only the target node (node B) responds with a name reply packet. The name reply contains the IP and station addresses and service information (supported NS services) in the path report. The other nodes simply discard the request. If there is no response from the target node, then a Probe proxy request is sent out

over the secondary multicast address. See "Probe Proxy Request/Reply" below for more information.

Node B responds to the name request with a path report that contains the IP and link-level station (MAC) addresses of the target node. The station address in the response received from the target node will only be considered valid if the IP or network numbers are the same as the originator's. If they are not, then the originator must still resolve the IP-to-station-address mapping. It does this with a Probe Virtual Address Request (VNA), as described in the next section. If a proxy server replies to the request, it may fill in the station address field with a null address or all zeros to ensure the requester generates a VNA for the correct station address.

Probe Virtual Address Request/Reply - VNA

This is the equivalent to ARP in the NS world. It provides for the mapping of a network (IP) address to link-level station address. The router will generate a VNA when it needs to forward IP packets to an NS node and there is no entry for it in its ARP cache. The router will have the destination IP address and requires the station address in order to forward the packet. The target node responds to the request with a path report that contains its station address. The router then forwards the IP packet onto the destination NS node. The router will respond to nodes generating these requests for destination networks known to it.

Probe Proxy Request/Reply

The proxy request packet is basically the same as the name request (see "Probe Name Request/Reply"), except that it is sent out over the secondary multicast address. The secondary multicast address is listened to only by Probe proxy nodes or servers. Proxy servers contain a database of name-to-IP-address mappings for some or all of the nodes in the network. The server will answer the requests for nodes that are on different subnets, or for local nodes that cannot respond to name requests, with the IP address from the database.

HP Router Probe Implementation

When HP routers boot, they transmit on all IP network interfaces the unsolicited reply announcing their presence. The routers use Probe VNA like ARP in order to obtain the station address of a destination node, and will try both Ethernet and IEEE 802.3 encapsulation to contact the target node. The router will answer VNA requests, for networks known to it, with its station address. The requester will then direct its IP traffic to the router. The routers do not implement name or proxy packets (requests or replies). This is not a problem if bridging is enabled on the same ports that enable IP routing in the network. The router will bridge these multicast requests to their target nodes and back to the requester. If bridging is not enabled, then other provisions must be made to resolve the name-to-IP mapping, by using proxy servers or statically building the Nodal Registry in each node.

Connection Scenarios

The following are two different TCP connection scenarios between node A and node B over a router network. The first is without a proxy server, with bridging enabled on the routers, and the second is with a proxy server and with bridging disabled on the routers.

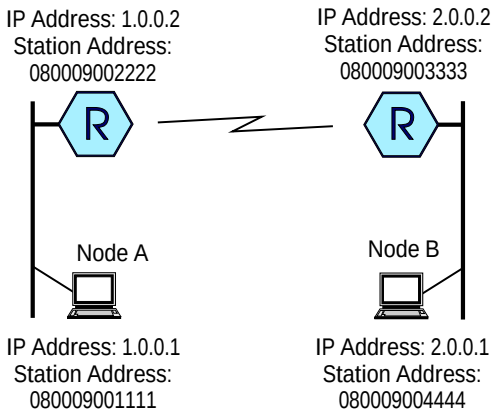

Figure 1. Network Diagram for Scenario 1

Scenario 1

Figure 1 shows the first connection scenario, without a proxy server.

- Both routers are configured for IP routing and bridging
- Node A wants to set up a TCP connection to node B, which is on the other side of the network.
- Probe is enabled on the routers.

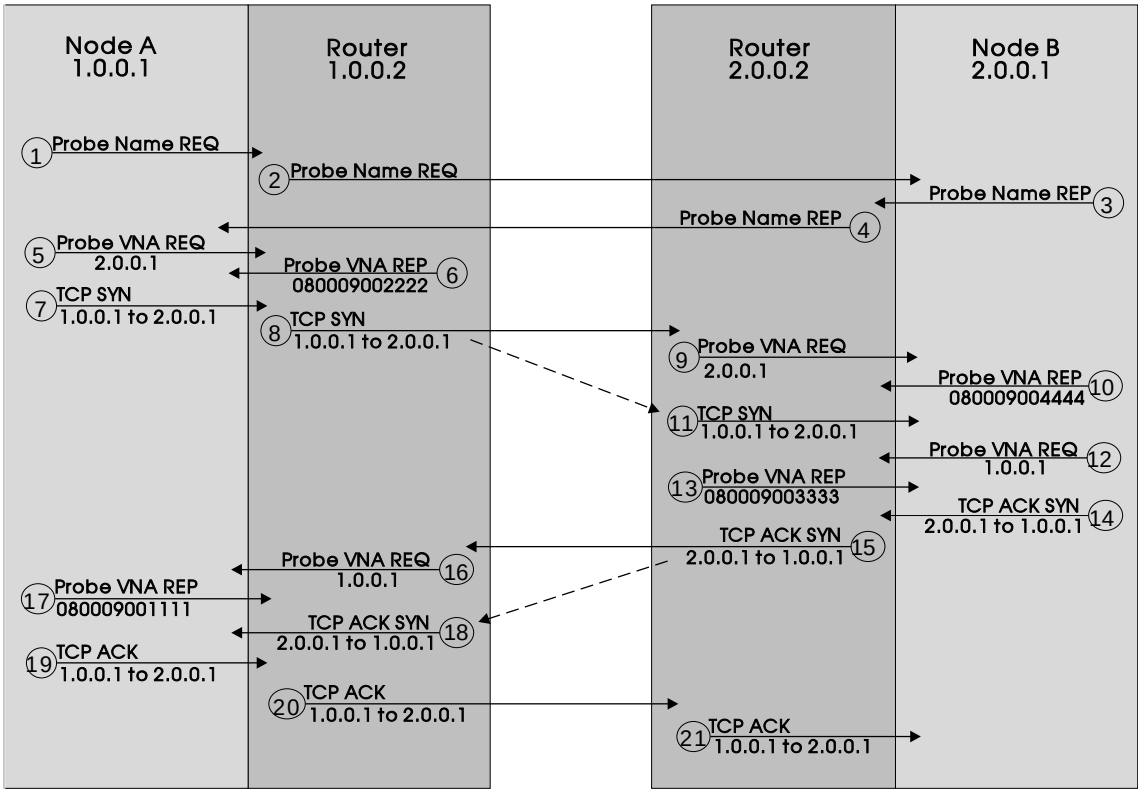

Figure 2. Connection Establishment Steps for Scenario 1

The steps required to establish the connection are shown in figure 2 above and are described below:

1. Node A wants to connect to node B, checks the Nodal Registry, and does not have an entry for node B. It sends out a Probe name request.

2. The name request from node A is bridged by the routers to node B.

3. Node B replies to the name request.

4. The name reply (containing the IP address) is bridged back to node A.

5. Node A then sends a VNA request using node B's IP address.

6. The local router knows a route to node B's network, so it responds to node A's VNA with the router's station address.

7. node A transmits a TCP sync packet to node B's IP address to the router's station address.

8. The local router routes the TCP sync packet to the destination network.

9. The remote router receives the TCP synchronous packet for node B. If node B is not in the ARP cache, it sends a VNA request (and ARP, if it is enabled) for node B's station address.

10. Node B replies to the router with a VNA with its station address.

11. The remote router sends the TCP sync packet to node B's station address.

12. When node B wants to reply to the TCP sync packet with a TCP ack sync packet, it must get the return station address, so it sends a VNA request.

13. The remote router responds to the VNA request with its station address.

14. Node B sends a TCP ack sync packet to node A's IP address to the station address of the remote router.

15. The remote router routes the TCP ack sync packet to the local router.

16. The local router receives the TCP ack sync packet for node A; if it is not in the ARP cache, it sends a VNA request (and ARP, if it is enabled) for node A's station address.

17. Node A replies to the router with a VNA with its station address.

18. The local router sends the TCP ack sync packet to node A's station address.

19. Node A does a lookup of the station address and transmits a TCP ack packet to node B's IP address to the router's station address.

20. The TCP ack is routed to the remote router.

21. The remote router checks its ARP cache then sends the TCP ack to node B's station address.

Scenario 2

A proxy server is used in this scenario.

- Each router is configured for IP routing only.
- Node A wants to set up a TCP connection to node B, which is on the other side of the network.
- Probe is enabled on the routers.
- There is a Probe proxy server on each subnet.

The steps required to establish the connection are as follows:

1. Node A wants to connect to node B, checks the Nodal Registry, and does not have an entry for node B. It sends out a Probe name request.

2. The name request is dropped by the local router.

3. After no response, node A sends out a proxy request.

4. The proxy server responds with the IP address of node B.

5. Node A then sends a VNA request using node B's IP address.

Steps 6 through 21 are the same as those in scenario 1.

Summary

In order to run NS on an IP router network, the main issue is how the name-to-IP-address mapping is done for nodes on other subnets. If bridging is enabled, then the name requests are bridged and the mapping is performed using Probe name request/reply packets. If bridging is not enabled on the network, then proxy servers need to be on each subnet, and the mapping is performed by Probe proxy request/reply packets. As an alternative, the Nodal Registry on each node communicating over the network could be configured with the entries required for nodes on remote subnets.

Data Compression for WAN Links

Technical Brief

Using data compression on wide area links has ramifications that must be considered when implementing routed network solutions. This note will do the following:

- Describe current methods of data compression.
- Briefly explain HP's compression algorithm.
- Describe the performance test method used for testing compression with HP routers.
- Report compression test results.
- Give some general guidelines for effectively designing networks that utilize compression.

What is Compression?

Compression is a method of encoding data such that the resulting file or packet is smaller than the original. Therefore, when the data is sent, say over a wide area link, it will utilize less bandwidth than the original data uses. At the other end of the link, the data is then uncompressed to its original size. In this way, one can gain improved WAN throughput without having to upgrade to higher speed lines. For example, using the HP AdvanceStack Router 650, running compression on 64 Kbit/s links can yield throughputs of up to 96 Kbit/s. See the "Design Guidelines" section for more details.

Data compression is achieved by replacing repeating strings of data with a smaller string called a "key". Compression algorithms use what is called a "dictionary", which contains the mappings of strings to their respective keys. The sending device removes the repeating strings within each packet and replaces them with the key string, and the receiving device reverses the process, restoring the original data and removing the key. Both the sending

and receiving devices must have the same dictionary in order to replace the key with the original text at the receiving end.

Running Dictionary

One method of data compression uses a “running dictionary”, meaning that the mappings of strings to keys is maintained and reused across multiple packets being transmitted and received. This method has the benefit of achieving high compression ratios, but it has some restrictions. First, running dictionaries require large amounts of memory to be maintained. Generally, the larger the dictionary (i.e., the more memory used), the better the compression results, because it is more likely that a string will already have a key defined in the dictionary. Another more serious drawback is that if the dictionaries at the sending and receiving devices gets out of synchronization, all packets will be dropped until the dictionaries re-synchronize. Therefore, if the link between the sending and receiving devices is not reliable, then it is highly probable that the dictionaries will often be out of synchronization, which may result in poor performance or complete failure of the connection. Running dictionaries have difficulty with WAN technologies that use datagram oriented-protocols (e.g., X.25 or frame relay), in which packets may arrive at the receiving device out of order or corrupted. In other words, running dictionaries require that packets be reliably delivered across the link with no data loss or corruption. Additionally, if the WAN link uses multiple virtual circuits, then a separate dictionary is used for each virtual circuit, thus requiring more memory.

Packet-by-Packet Dictionary

To overcome the memory and link quality limitations of running dictionaries, another method of data compression was developed in which a very small dictionary is used, and it is reset for each packet. Therefore, very little memory is required to store the dictionary, and there is no way for the sending and receiving devices to get out of synchronization, even when using unreliable WAN link protocols. Packet-by-packet dictionary methods do not achieve compression ratios as high as running dictionary methods, because there is no chance of reusing a key from matching strings in previous packets. However, the higher ratios of running dictionary methods are diminished by the fact that reliable link protocols use more link bandwidth for acknowledgment packets. Therefore, the overall throughput may not be as high as expected when using running dictionaries. Additionally, the synchronization problems associated with running dictionaries makes the overall performance of packet-by-packet methods more effective for use on wide area links.

HP's Compression Algorithm

Beginning November 1, 1993, HP has been shipping a variation of the Lempel-Ziv (LZ) lossless compression algorithm on all HP routers. The HP scheme (called HP Packet-by-Packet Compression, or HP PPC) compresses each packet independently using the packet-by-packet dictionary method in which the dictionary is reset with each packet. Also, HP PPC ensures that the compressed packet is never larger than the original, otherwise the original (uncompressed) packet will be sent. Sending the uncompressed packet in this case saves CPU bandwidth by eliminating the need to uncompress the packet at the other end of the link. Additionally, sending the uncompressed packet saves WAN link bandwidth by assuring that a packet is never expanded beyond its original size. Not all compression algorithms guarantee smaller, or even equal size data!

Run Length Encoding

In addition to packet-by-packet compression, the HP PPC algorithm employs an additional encoding scheme, called run length encoding, to further compress data. Run length encoding is a method of replacing repeated occurrences of a certain character with a single occurrence, followed by the number of times it occurs (i.e., the run length).

Currently, HP PPC is supported only on HP point-to-point links. Later releases of router software will support other, more standard link technologies, such as Point-to-Point Protocol (PPP), including PPP over ISDN (V.25 bis) and PPP over frame relay. HP PPC uses very little memory, and is optimized for running on WAN links that use non-reliable, datagram-oriented protocols, such as HP point-to-point links running the LLC1 datagram service. You can also use the LLC2 (reliable) service, which provides error detection as well as error recovery by retransmission. However, LLC2 uses more link overhead in acknowledgments and retransmissions, and is not required for HP PPC to run efficiently.

Compression Performance Testing

HP conducted performance testing using the Calgary Corpus test files, which are industry-standard files for performance testing using various types of data. The table below describes each file in the Calgary Corpus set.

File name	File Description	File Size (bytes)
bib	Bibliographic files	111261
book1	Hardy: Far form the Madding Crowd	768771
book2	Witten: Principles of Computer Speech	610856
geo	Geophysical data	102400
news	News batch file	377109
obj1	Compiled code for VAX: compilation of progp	21504
obj2	Compiled code for Apple Macintosh: knowledge support system	246814
paper1	Witten, Neal and Cleary: Arithmetic coding for data compression	53161
paper2	Witten: Computer (in) security	82199
paper3	Witten: In Search of "Autonomy"	46526
paper4	Cleary: Programming by Example Revised	13286
paper5	Cleary: A Logical Implementation of Arithmetic	11954
paper6	Cleary: Compact Hash Tables Using Bidirectional Linear Probing	38105
*pic	Picture number 5 from the CCITT facsimile test files (test + srawings)	513216
progc	C source code: compress version 4.0	39611
progl	LISP source code: system software	71646
progp	Pascal source code: prediction by partial matching evaluation program	49379
trans	Transcript of a session on a terminal	93695

* The pic file was left out of the testing because it contains mostly zeros, which skewed the compression results in favor of higher rations. For example, just copying the pic file by itself yielded approximately 5:1 compression ratio. By leaving out the pic file, we achieved a file set that more closely approximated test files.

Using these test files, HP routers were set up in various configurations and file transfers were done across HP point-to-point links to determine acceptable performance and throughput. Tests were conducted using both IP and IPX protocols. The following section gives the results of this testing, and compares them with results reported on other popular compression algorithms. The test results lead to some guidelines for employing compression on WAN links in routed networks. Those guidelines are detailed in the "Design Guidelines" section following the test results.

Test Results

Using the Calgary Corpus files, observed compression ratios averaged 1.5:1. Beware of vendors who claim very high compression ratios, as ratios are entirely dependent on the content of the data being transferred. Basically, any vendor can "customize" compressible text and claim ratios based on those numbers. The reality is that with random LAN traffic, using packet-by-packet compression techniques over non-reliable wide area links, 1.5:1 is a reasonable number.

Design Guidelines

As a result of testing, the following guidelines should be observed when implementing compression on HP Point-to-Point links.

1. **HP Router 650 (card cage model holding up to 4 interface cards or 16 ports):**

- **Compression should not be run on WAN links with speeds greater than 256 Kbit/s.**
  Since the HP PPC algorithm is implemented in software, it does require CPU bandwidth to run. If compression is run on higher-speed links, the throughput will likely not increase beyond that of a 256 Kbit/s link due to the required CPU bandwidth to compress and decompress the data. This fact is true with any software-based compression algorithm. The only way to achieve adequate compression ratios on higher-speed links is to use a more expensive, hardware-based compression algorithm. As a general rule of thumb, the aggregate throughput numbers given below can help you determine how many and what speeds of links will run successfully on the Router 650.

- **Each 4-port synchronous interface card has an aggregate throughput of 306 Kbit/s to 460 Kbit/s with compression running.**
  The Router 650 is designed such that each interface card has its own processor, so all compression processing is done on the interface card. There were no negative effects on performance when four 64 Kbit/s links per card were run with compression; in fact, even fully loaded (with bi-directional file transfers), each link maintained a throughput of 96 Kbit/s, a 50% increase in bandwidth utilization when compared to sending the same data over the same links without compression running. See the discussion on aggregate throughput, below, for more details.

Aggregate Throughput

Generally speaking, the CPU is shared among the four WAN links on the card. Therefore, for each interface card, adding all the link speeds will determine that card's aggregate throughput.

If the total of all link speeds for a given card is less than 306 Kbit/s per 4-port synchronous card, then all links will be able to handle maximum load with no problems. This 306 Kbit/s value is very conservative, as it takes into consideration the maximum amount of traffic going in both directions on the link simultaneously.

If the total of all link speeds for the card is between 306 Kbit/s and 460 Kbit/s, then the compression algorithm may not keep up with all links running at maximum capacity. That is, throughput may not increase as link speed increases. However, normal day-to-day activities (non-peak loads) should still perform with no problems.

If the total of all link speeds for the card is greater than 460 Kbit/s, then turn off compression or reduce the link speed on one or more links until the total is less than 460 Kbit/s.

For example, using four 64 Kbit/s links would yield a total of 256 Kbit/s. Since the total is less than 306 Kbit/s, there will be no performance limitation. Similarly, one 256 Kbit/s link per card can run with no limitations.

WAN link planning

When planning a wide area network, it may be useful to know what link speeds are needed given an estimated WAN utilization level. For example, if you know you will utilize approximately 84 Kbit/s of WAN bandwidth, then what link speed should you purchase for running compression? You can use the following formula to determine the desired link speed:

(throughput) / 1.5 = (link speed)

Therefore, given a desired throughput of 84 Kbit/s, you may purchase a WAN service with a link speed equal to (84) / (1.5), or 56 Kbit/s.

2. Fixed-port routers (HP Router ER/TR/SR/FR/PR/FR/PR/TFR etc.)

- **Compression should not be run on WAN links with speeds greater than 64 Kbit/s.**
  As stated above, in the Router 650 case, if compression is run on higher speed links, the throughput will likely not increase beyond that of a 64 Kbit/s link due to the required CPU bandwidth to compress and decompress the data. The 64 Kbit/s limit is imposed on the these routers, instead of the 256 Kbit/s limit of the Router 650, because the architecture of fixed-port routers is such that there is one processor for both routing and compression; thus, CPU demands are greater.
- **On an HP Router SR, no more than 2 WAN links should be configured for compression and the third WAN link should be configured for speeds no higher than 9.6 Kbit/s.**
  Again, the software-based compression algorithm will prevent higher throughputs from being achieved. However, if the two links configured for compression are not completely saturated, then it will be entirely possible to get more than 9.6 Kbit/s throughput on the third link.

3. For all HP routers:

- **When running slow WAN links (9.6 Kbit/s to 19.2 Kbit/s) it is usually beneficial to run compression.**
  Throughput will generally be improved when running compression on very slow links, so it is a good idea to turn on compression for these links except where noted below.

- **In IPX environments, run compression only under the following conditions:**

  File server: Must be running the packet burst module, **PBURST.NLM version 2.02** or greater. Also, the following parameter must be set in the STARTUP.NCF file:

  set maximum physical receive packet size = 1518

  Client station: Must be running **VLM 1.1** or greater, and the following parameters are recommended for the NET.CFG file:

  MINIMUM TIME TO NET = 0 (default)
  PBURST READ WINDOW SIZE = 16 (default)
  PBURST WRITE WINDOW SIZE = 10 (default)

  Higher values were used for the window sizes, but there was little performance difference. Also, increasing the window size could result in many more retransmissions because if a packet is lost, all packets in the window must be retransmitted, not just the one packet that was lost.

  The reason for the burst mode requirement in Novell networks is that the standard NCP (NetWare Core Protocol) uses a method that requires the acknowledgment of every packet that is transmitted. Additionally, standard Novell file servers acting as IPX routers require that data packets be no larger than 512 bytes. Considering that the IEEE 802.3/Ethernet standards allow for a maximum packet size of 1514 bytes (and 4500 bytes in token ring networks!), one can see that even in non-compression implementations, valuable bandwidth is being wasted. It is for this reason that Novell implemented PBURST.NLM. PBURST not only allows for larger packets, but also allows bursts of up to 64 kilobytes of data to be sent before an acknowledgment is required. Performance testing has shown that without packet burst mode *and large packets*, compression is not useful, and may actually *decrease* performance.

Versions of PBURST.NLM earlier than 2.02 had problems that also led to insufficient performance in both compression and non-compression environments. It is therefore recommended that version 2.02 or later be used. Similar reasoning is used in the VLM 1.1 client software recommendation.

Conclusion

There are many tradeoffs associated with the various compression techniques available on the market. Particularly, with software-based algorithms, understanding the effects on CPU utilization becomes critical to designing router networks that maximize available bandwidth. Performance testing has shown that there are no performance issues with the HP AdvanceStack Router 650 running compression on all WAN links at 64 Kbit/s, or combinations of compression and non-compression links up to 256 Kbit/s. With the fixed-port routers, there are certain limitations of the software-based HP PPC algorithm that impose some design constraints with respect to supported link speeds and number of links that can run compression on a given router. Compression ratios can vary widely depending on the type of data used for the compression test; however, the HP PPC ratios are approximately 1.5:1 for random LAN traffic.

3

Application Notes and Case Studies

Application Notes and Case Studies

- Improving Network Availability
- ISDN Wide Area Network Design:
  Dry Creek Joint Elementary School District
- Shining a Light on FDDI
- Using Synchronous Pass-Through
  to Consolidate Synchronous Traffic
- Routing with OSPF
- Linking Up withFrame Relay
- Frame Relay Network Design:
  Fleet Call, Inc.,

Case Study

Improving Network Availability

Updated 7/93

For companies that rely heavily on network applications, a failure in the network can be disastrous. Network failures are more likely to occur as networks grow—the natural result of employing more and more network equipment. This application note examines some of the methods for increasing network availability (uptime) in router-based networks and thereby reducing or eliminating user-perceived network failures.

The most common technique for increasing network availability is to provide alternate paths for data in the event of link or router failures. This plays to the strength of routers—quickly detecting network failures and routing data around them. This application note examines the different ways of providing alternate data paths to survive both data-link and router failures.

Permanent Alternate Paths

Consider the network in figure 1. There is only one path for data from any site to any other site. Should link 1 fail, for instance, systems at site A will be unable to communicate with systems at either site B or site C, until link 1 has been restored. Adding a third link, as indicated by the dotted line between the site A and site C routers, provides an alternate path for data to and from sites A and C. Once the third link has been installed, any single link failure will not result in a user-perceived network failure.

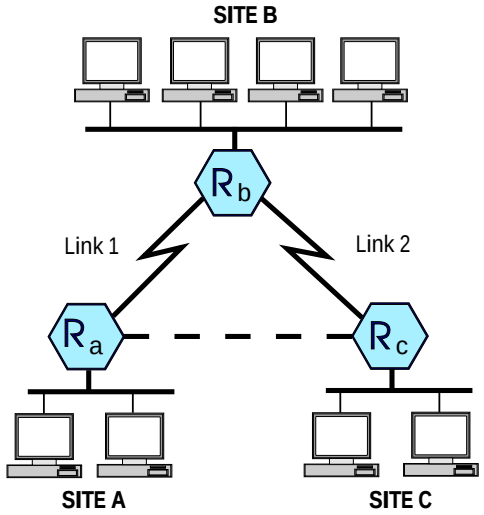

Figure 1. Improving Availability with an Alternate Path

Data is routed between any two networks on the lowest-cost path. The lowest-cost path from a router to a destination network is determined by the router's redirectors. Redirectors exist for each of the different network protocols—IP, Novell IPX, AppleTalk, DECnet, XNS, and the learning bridge. Redirectors react to network changes by sending, receiving, and processing routing updates. They calculate a forwarding table from the available routing information. The forwarding table indicates the cost to the destination network and the interface (circuit group) on which data bound for the destination network will be transmitted by the forwarder (the process within each redirector that routes data packets).

During normal router operation, the forwarding table indicates only the best (lowest- cost) path to a destination network. When link 1 fails, router A and router B will exchange routing information with router C. Router A will learn about its new path to the site B LAN (through router C), and router B will learn about its new path to the site A LAN (again through router C). Note that it takes a short period of time for the routers to detect the link failure and discover new routes to all available networks. This time is referred to as convergence time. Small networks will generally converge within one minute.

The permanent alternate path added in figure 1 has several advantages:

- *Immediate availability*. It can be used (to route data from site A to site B) as soon as the routers determine that it is the best path to a network that is unreachable due to another link failure. There are no

setup delays or other problems that can occur with switched backup schemes.

- *Flexible networking options.* Although the permanent alternate path selected for wide-area networks is typically a leased line, it could be a public switched service such as X.25, frame relay, or SMDS. In addition, there is a wide range of link speeds from which to choose.
- *Performance.* Better throughput and response times for traffic (if any) between sites A and C (in figure 1) are easily achieved.

Coterminus Circuits

Another way to establish an alternate path is to create a circuit group with two or more coterminous circuits. A circuit group is a logical structure that contains one or more circuits. Coterminous circuits are simply circuits that have common starting points and common end points.

Consider the network in figure 2. Two coterminous circuits are used to connect routers A and B. Both circuits have been configured as members of a single circuit group on each router. The forwarding process within each redirector directs traffic to circuit groups. The forwarders are not aware of the lower circuit-level detail. Thus, each router in figure 2 has one route to the LAN attached to the other router. However, data is routed on both circuits in a load-sharing fashion.

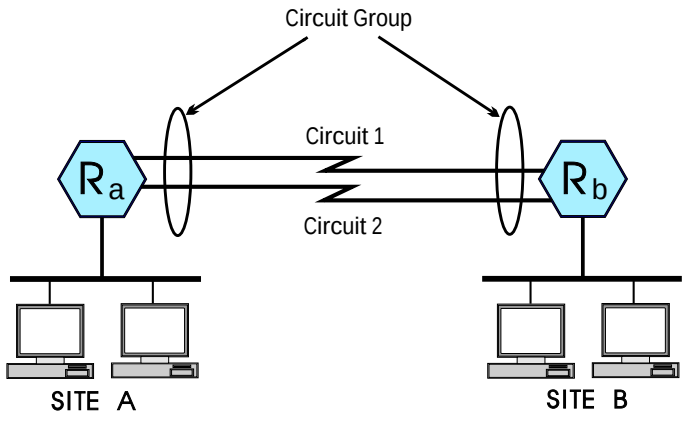

Figure 2. Interconnecting Routers with Coterminus Circuits

If one circuit fails, traffic is carried on the remaining circuit(s). As long as one circuit within a circuit group is functioning, the redirectors are not informed of circuit failures and thus do not try to find an alternate route. When the failed circuit is restored, it will automatically be used for sending traffic to the remote router. Since a circuit group did not fail, an alternate route is not selected. Instead, traffic is aggregated on the remaining circuit(s) within the group. Thus, convergence is not an issue with this type of circuit failure. The time to detect and remove a failed circuit is typically 2 to 20 seconds, depending on the characteristics of the failure.

For bridging, must the Spanning Tree Protocol be enabled to eliminate the loop created by the additional load-sharing circuit? No, the Spanning Tree Protocol is not required since logically there is no loop. The bridging redirector directs traffic to the single circuit group, rather than the two circuits.

There are limitations on the use of coterminus circuits. To understand those limitations you need to know how the circuit group manager works.

Circuit Group Manager

After a routing or bridging forwarder has received a packet and has determined that the packet must be forwarded, it hands the packet to the circuit group manager with instructions to send it on a particular circuit group. When multiple circuits exist in a circuit group, the circuit group manager decides which circuit to use based on its circuit assignment algorithm. One of two basic algorithms is used, depending on the type of the packet.

Circuit Assignment Algorithm #1: Random

The first algorithm assigns traffic on circuits randomly. That is, when a packet is being transmitted on a circuit group with two circuits, the choice of which circuit to use is made on a random basis. Table 1 shows the packet types for which the the random circuit assignment method is used.

Table 1. Packet Types for Random Circuit Assignment

Protocol	Packet Type
AppleTalk	Zone Information Protocol (ZIP)
AppleTalk	Routing Table Maintenance Protocol (RTMP)
AppleTalk	Name Binding Protocol (NBP)
AppleTalk	Address Resolution Protocol (AARP)
AppleTalk	Echo Protocol (AEP)
DECnet	Routing Protocol
Learning Bridge	multicast/broadcast packets
Novell IPX	all packet types
XNS	all packet types

Circuit Assignment Algorithm #2: Indexed

The second circuit assignment algorithm is referred to as “indexed”. The indexed algorithm typically MODs the sum of the source and destination addresses by the number of available circuits in the circuit group. Recall that the MOD function returns the remainder of a division. The following example illustrates this algorithm.

The indexed circuit assignment algorithm is used for all IP data packets. In figure 2 (above), assume that a computer at site A with IP address 15.8.128.210 sent a packet to a computer at site B with IP address 15.8.64.173. To determine which of the two circuits to use for sending the packet, the circuit group manager adds the two IP addresses. The result in hexadecimal is 1E10C17F. This value is then MODed by 2 (the number of available circuits in the circuit group). The result is 1 (the remainder when 1E10C17F is divided by 2). A MOD value of 0 indicates that the packet must be transmitted on circuit 1 and the MOD value of 1 indicates that the packet must be transmitted on circuit 2.

Suppose one of the coterminus circuits fail. What happens? Again, consider the network depicted in figure 2. When a packet with the IP addresses given above is received by router A, the two addresses are MODed by 1 instead of 2. This causes the packet to be transmitted on the one remaining circuit. When the failed circuit is restored, the indexed algorithm will operate as before.

Table 2 shows the packet types for which the indexed circuit assignment algorithm is used. Also shown is the index used for the calculation.

Table 2. Packet Types for Indexed Circuit Assignment

Protocol	Packet Type	Index
AppleTalk	data packets	last 4 bytes of destination station address + last 4 bytes of source station address
DECnet	data packets	destination node # (intra-area traffic) or destination area # (inter-area traffic)
IP	all packets	destination IP address + source IP address
learning bridge	no filters or filters without load balancing	last 4 bytes of destination station address + last 4 bytes of source station address
learning bridge	load balancing filters	none; traffic sent by type on assigned circuit

Load Balancing Overrides

The circuit assignment algorithms described above can be overridden for IP and Novell IPX traffic. Answering Yes for the IP network interface definition's Load Balancing parameter changes load balancing for all IP packets on that particular circuit group to random (circuit assignment algorithm #1). Random usually works quite well and is generally better than the indexed method when the number of simultaneous sessions (source/destination address pairs) is low.

Answering No for the Random Load Balancing parameter in the Novell IPX network interface definition selects indexed (circuit assignment algorithm #2) load balancing rather than random load balancing on that particular circuit group. Indexed load balancing is recommended when systems running Netware have been configured to use Netware's Burstmode NLM.

Load Balancing Limitations

First, the circuit assignment algorithms ignore such obvious parameters as delay and throughput. This means that circuits placed in a circuit group must be of the same capacity. Otherwise, low-capacity circuits (56 Kbit/s) will have to handle the same networking load as high-capacity circuits (1.544 Mbit/s).

Second, the circuit assignment algorithms do not necessarily provide balanced link utilization. Large packets may be sent on circuit 1 while small packets may be sent on circuit 2. Additionally, the indexed circuit-assignment algorithm relies on a large enough population of users (actually, summed addresses) to achieve a measure of balanced circuit utilization.

As a final note, it is generally a good practice when designing a network with coterminous circuits to purchase the circuits from different carriers to insulate your network from multiple simultaneous circuit failures.

Packet-Switching Networks

Packet-switching networks (PSNs) based on technologies that include X.25, frame relay, and Switched Multimegabit Data Service (SMDS) are designed to be very reliable. Reliability is achieved through a highly redundant network architecture with many interconnected, parallel switching nodes; see figure 3. The relatively recent deployment of frame relay and SMDS, which offer high throughput and low latency, have added to the desirability of these networks for interconnecting LANs.

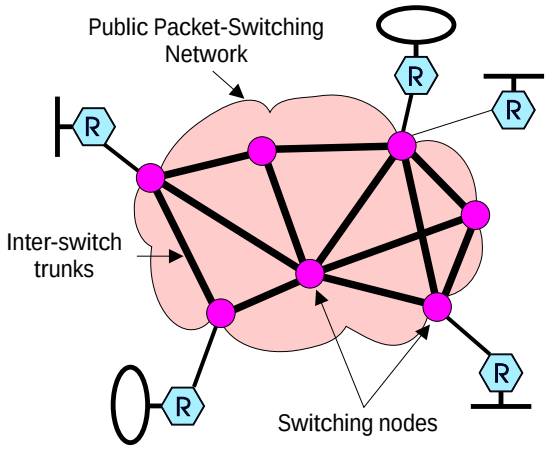

Figure 3. LANs Connected through a Packet-Switching Network

Although packet-switching networks are themselves typically highly reliable, the access circuit that attaches a router to the network is subject to failure and thus is an obvious point of vulnerability.

The switching nodes used to construct PSNs usually have fault-tolerant features. However, connecting multiple routers to the same switch increases the seriousness of a switching node failure if it does occur. See the two routers in the upper right in figure 3.

Finally, when choosing a PSN to interconnect routers, ask your network vendor for a description of the physical network to verify that the network has the redundancy you expect.

Circuit-Switching Networks

Circuit-switching networks are those in which a physical circuit is established after a number is dialed. Circuit-switching networks include ISDN, Switched 56 (Kbit/s), Switched 384 (Kbit/s), and the Public Switched Telephone Network (PSTN).

Circuit switching is a fundamentally different concept from packet switching. In a circuit-switching network the subscribers or users must dial a phone number to access a remote device such as a router. When the number is dialed, a data path is established by the circuit-switching network. Once this data path is established, communications may proceed between the two connected devices. Users are typically charged for service on the basis of connect time. Circuit-switched communications are thus point-to-point and limited in duration. In this application note, "switched circuit" is used as a generic term referring to connections established over circuit-switching networks.

In packet-switching networks, the data path between subscribers or users is permanently established. Users communicate with each other using data-link-layer or network-layer addresses. The data path between any two users in a packet-switching network is usually referred to as a virtual circuit. Users are typically charged for service on the basis of the number of packets transmitted. Packet-switched communications are thus multi-point and may or may not be limited in duration.

Circuit Group Considerations

The way in which a switched circuit operates is governed primarily by its type of circuit group membership. Switched circuits can be configured as:

- Primary circuits.
- Backup circuits.
- Pool circuits.

Note that the circuit-switching capabilities discussed in this application note pertain only to the HP Router ER, FR, SR, and TR. In particular, these capabilities are *not* available on the HP Router CR.

Primary Circuits

A primary circuit is configured when a circuit is added to a circuit group as a "circuit group member". This is how all circuits other than switched circuits are usually configured. Switched circuits can be configured as primary circuits. This is generally the way remote sites that connect to central site pool circuits are configured; see figure 4.

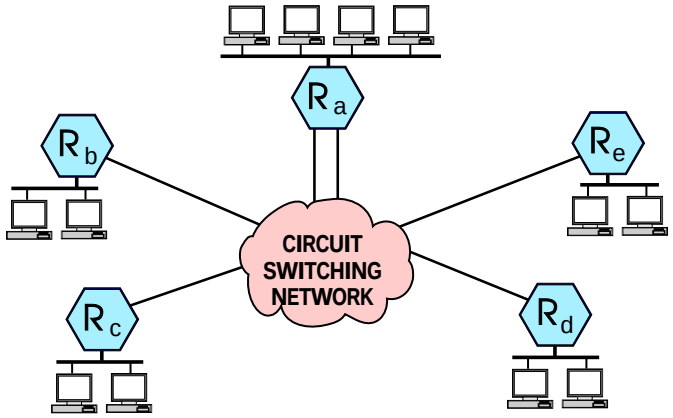

Figure 4. Routing Using Switched Network Services

Backup Circuits

Switched circuits are often the most cost-effective alternative to a meshed network (a network using permanent alternate paths) for improving availability. Switched services may be used to back up private leased lines or packet-switching networks. When compared to leased-line CSUs/DSUs with dial backup capability, the HP routers' switched backup function provides more flexibility in selecting the type of switched circuit.

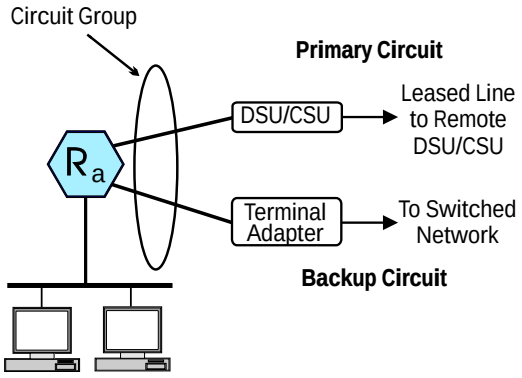

Figure 5. Backup and Primary Circuits in a Circuit Group

A backup circuit is defined when the circuit is added to a circuit group as a "backup circuit group member". Figure 5 shows a circuit group with a primary and a backup circuit. A backup circuit is enabled when the primary circuit(s) in a circuit group fails. The backup circuit is disabled when a primary circuit is restored. Traffic is transmitted on either the primary circuit or the backup circuit but not on both at once. Thus, backup circuits cannot be used for bandwidth-on-demand applications. Additionally, dynamic routing protocols are normally used on backup circuits since they are operational only when the primary circuit has failed.

Primary and backup circuits are assigned to the same circuit group and thus share network interface definition attributes such as addresses. Backup circuits are recommended for use with all of the routable protocols. Since a backup circuit is normally inactive, a switched (dial-up) circuit is the logical choice to implement as a backup circuit, since it is the most cost-effective alternative.

When a switch from primary to backup circuits must be performed quickly, ISDN and Switched 56 are preferred to analog modems. ISDN and Switched 56 connections can typically be completed in 2 to 5 seconds, versus 25 to 30 seconds for analog modems.

Pool Circuits

Pool circuits (circuit group pool members) can be used as primary or backup circuits. All pool members of a circuit group have the same attributes of the network interface definition, since they are members of the same circuit group. Thus, they share the address defined for the circuit group. IP pool circuits can be used only to transmit and receive IP traffic. Figure 6 shows a circuit group with several pool circuits.

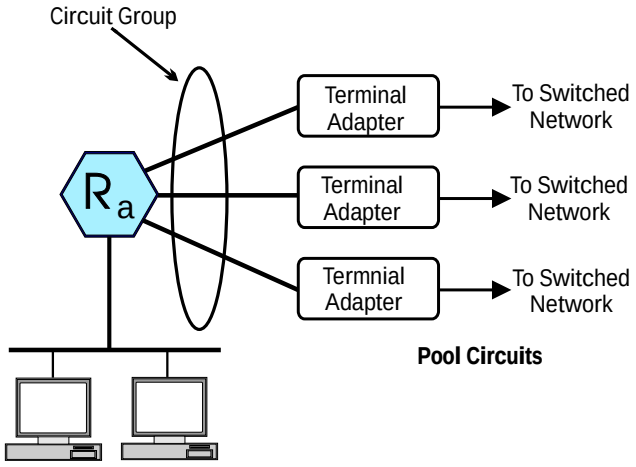

Figure 6. Pool Members of a Circuit Group

The use of pool circuits is ideal in a couple of situations. The first situation is when the traffic volumes are low and transmissions are infrequent. Transmissions at the end of the business day to update databases are an example. In this scenario, several (or many) sites are networked together using switched services. Figure 4 (above) shows a network utilizing pool circuits for this purpose. The second situation is the use of pool circuits to provide a backup path for failures of leased lines or packet-switching networks. Figures 7 and 8 illustrate these cases.

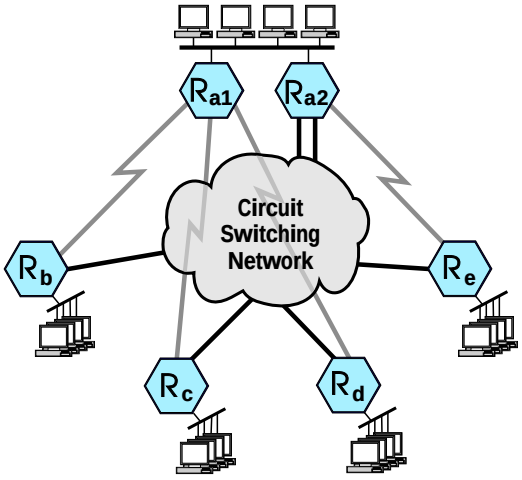

Figure 7. Leased-Line Network with Circuit-Switching Backup

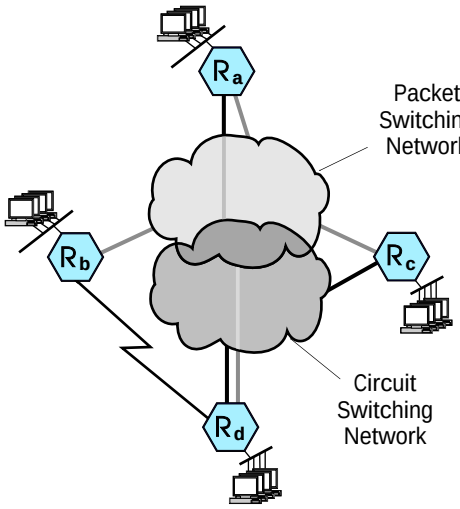

Figure 8. Packet-Switching Network with Circuit-Switching Backup

Minimizing Connect-Time Costs

One of the most important issues with circuit switching is minimizing the cost of connect time. Several connection controls have been provided to manage the opening and closing of connections to avoid connection charges. These controls are described later in this document.

Since all dynamic routing protocols such as RIP generate periodic traffic to maintain routing tables, careful consideration must be given to the use of these protocols on switched circuits. Static routes can be used to effectively eliminate routing-protocol control traffic on switched circuits. This ensures that switched circuits are opened only for the transfer of user data.

Switched Circuit Types

The routers' Configuration Editor allows the definition of two types of switched circuits:

- Manual adapter
- V.25 bis adapter

These two circuit types can provide access to virtually any circuit-switching facility including ISDN, Switched 56 or 384, and the Public Switched Telephone Network. The following discussion assumes the use of an ISDN network, so the term "adapter" is used instead of DSU/CSU. However, note that the router functionality being discussed is independent of ISDN.

Manual Adapter

A manual adapter refers to any DCE (ISDN terminal adapter, modem, or DSU/CSU) that will initiate a connection to a remote DCE when the router raises the DTR lead (the data-terminal-ready signal) on the interface to the DCE. This is simply hardware-level signaling to initiate and break connections. Any of the synchronous interfaces (RS-232/V.24/V.28, V.35, RS-449/422/V.36, and X.21) may be used to connect to a manual adapter. Connections with a manual adapter may be initiated in one of two ways:

- When data is available
- When the circuit is enabled

There are four combinations of parameter settings to initiate connections. That is because there are two routers involved in setting up a switched connection, and there are two connection initiation methods on each router. Table 3 shows the four possible combinations. Of the four combinations, only the two shaded combinations are recommended. Thus, one side opens the circuit when data is available and the other side always waits to receive data.

Table 3. Parameter Combinations for Initiating Manual Adapter Connections

Side A Router	Side B Router
Circuit is enabled	Circuit is enabled
Circuit is enabled	Data is available
Data is available	Circuit is enabled
Data is available	Data is available

When the Connection Initiation parameter is set to “circuit is enabled” on both sides, the circuit will never hang up. The user will incur charges for a permanently open circuit!

When the Connection Initiation parameter is set to “data is available” on both sides, the circuit will never get established! HP routers do not implement the physical “ring indicator” signal, so there is no way to know about the arrival of an incoming call.

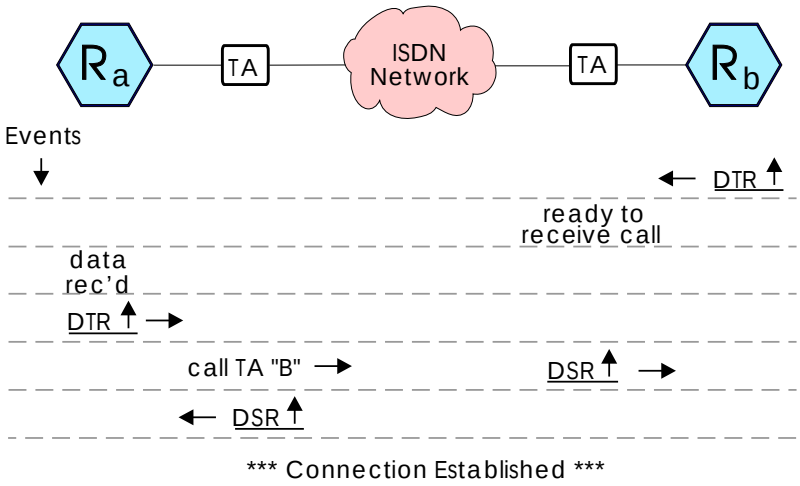

Figure 9. Connection Sequence for Manual Adapters

Figure 9 shows the events in the connection sequence for two routers using manual adapters. Router B is configured to connect when the circuit is enabled (during router bootup). Terminal adapter B has been configured for auto-answer (that is, not to dial a number when the router enables the circuit but rather to wait for an incoming call).

Router A is configured to connect when data is available. Data available means that the the router has received a packet which must be sent to router B. Terminal adapter A is configured to dial terminal adapter B's phone number when router A raises DTR (data terminal ready).

In this configuration, router A always initiates calls and router B always receives calls. In other words, applications on systems attached to router A always initiate the connections to systems attached to router B. This is frequently the situation for client systems at a sales/support location connecting to servers at a central or regional office for database access, credit verification, electronic mail, etc.

When circuit-switching network connections are used as backup circuits, and manual terminal adapters are used, the configuration just discussed is recommended.

When the establishment of connections must be bilateral, that is, when either router may initiate the connection establishment, a V.25 bis adapter must be used. Connecting with V.25 bis adapters is discussed further below.

Connection Controls—Manual Adapter

Using the connection method discussed in the previous section, several connection control parameters are useful in situations in which the destination is busy and in which connections fail due to problems in the network itself. These parameters include:

- Connect Retry Count
- Connect Wait Time
- Delay After Connect Failure

Two additional parameters can be used to minimize the connect-time charges. These include:

- Minimum Connect Duration
- Connect Inactivity Time

In many locations the telephone company or PTT bills the user one rate for an initial period of time and a second rate for incremental periods of time. For example, a telephone company in California bills a user at one rate for the first three minutes of a call and at a second rate for every additional minute. Minimum Connect Duration could be used, in this example, for the first three-minute period, and Connect Inactivity Time could be used for the incremental one-minute periods.

These parameters, when set to the correct billing periods, cause the router to monitor data transmission and reception. They serve two important functions:

- First, the user will incur the minimum connect-time charge for data transmissions. The router will automatically disconnect when data has not been sent or received for an incremental billing period.
- Second, after data has been sent or received, the connection is maintained until the end of the current billing increment. For example, suppose a connection is established and data is transferred for several seconds. Rather than immediately closing the circuit, the router holds the circuit open until the end (actually, one second before the end) of the Minimum Connect Duration. This prevents incurring the cost of another call if additional data is exchanged in this time period.

V.25 bis Adapter

V.25 bis is a CCITT data-communications standard that defines a set of commands exchanged between a DTE and a DCE. These commands control setting up and tearing down switched communications services. V.25 bis is similar in concept to the Hayes AT command set used in PC-to-modem communications. Unlike the AT commands, V.25 bis commands are used on a synchronous interface. Many ISDN terminal adapters and Switched 56 DSUs/CSUs now support V.25 bis. In addition, V.25 bis is available on some high-speed analog modems (when set up for synchronous operation).

HP routers support V.25 bis, plus extensions to V.25 bis that have been defined by Cisco Systems, Inc., and Ascend Communications, Inc. The V.25 bis extensions are primarily used to control parameters used with inverse multiplexers such as the Ascend Multiband.

The use of V.25 bis adapters provides a superset of the capabilities described previously for manual adapters. Additional capabilities provided include:

- Bilateral communication initiation and establishment. Each side of the switched connection can dial the other party and establish communication. V.25 bis supports connecting when “data is available or on incoming calls”, rather than when “data is available” as is the case using manual adapters.
- Enhanced security. A router may be set up to accept calls only from a predefined list of originating phone numbers (on adapters that return incoming numbers).
- A circuit-switched port can be configured to dial several different numbers. In the event that the first number dialed encounters a problem (such as a busy condition), another number can be dialed.
- Interface for parametric control of devices such as the Ascend Multiband inverse multiplexer for bandwidth-on-demand applications.
- Pooled circuit switching (IP only). This feature allows a circuit-switched port to receive or originate calls to a maximum of 16 different networks or subnetworks. Circuit-switched ports may be pooled such that up to three ports can be used as interfaces to a maximum of 16 (sub)networks.

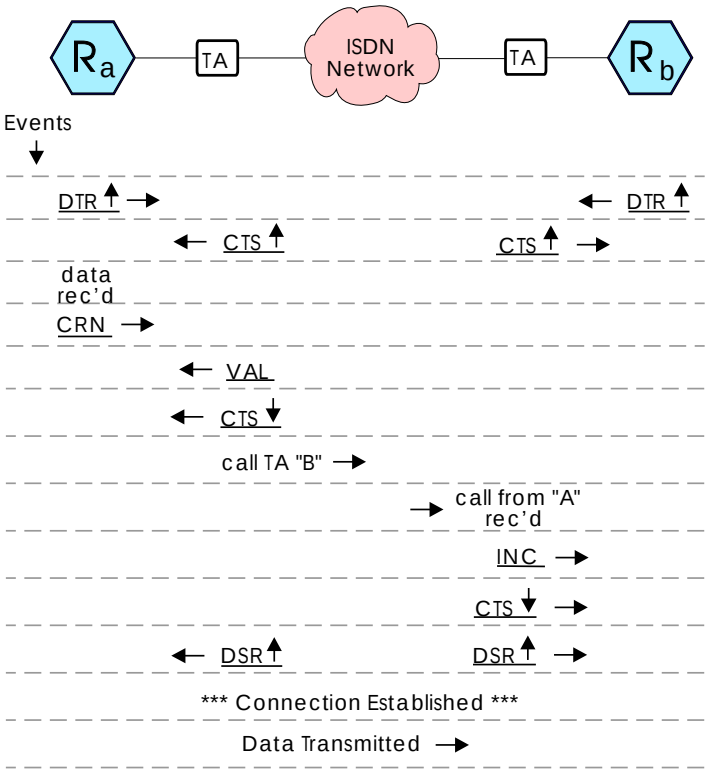

Figure 10. Connection Sequence for V.25 bis

Figure 10 shows the sequence of events in a successfully initiated connection between routers using V.25 bis terminal adapters. To signal that it is ready to engage in a V.25 bis dialog, the adapter raises its CTS lead (clear to send). The "CRN" sent by the router to the adapter is a connection request and the phone number to dial. The adapter responds with "VAL", which simply indicates that the request was formatted properly and that the adapter was in the proper state to receive the request. The adapter lowers CTS when an answer tone is detected to disable further V.25 bis dialog. The call is then placed.

On the receiving side, the adapter that previously raised CTS sends the router an "INC" to indicate that an incoming call has arrived. The originating phone number and the first part of the data packet is forwarded to the router; this assumes that ANI (automatic number identification) is available. Thus, the "ring indicator" control signal is not required for detection of an incoming call since that is obtained using the V.25 bis "INC" indication. If the

router will not accept the call (based on configured call restrictions), it immediately drops DTR. Otherwise the adapter subsequently drops CTS. As soon as the switched circuit is complete, DSR (data set ready) is raised by each adapter and the circuit is established.

Any of the synchronous interfaces except X.21 (that is, RS-232/V.24/V.28, V.35, RS-449/422/V.36) may be used to connect to a V.25 bis adapter. The X.21 interface lacks the necessary hardware-control signals required to implement V.25 bis.

Connections may be initiated in one of two ways with a V.25 bis adapter:

- When data is available or on incoming call
- When the circuit is enabled

Table 4. Valid Connection Initiation Parameters for V.25 bis Circuit Group Members and Backup Members

Side A Router	Side B Router
Circuit is enabled	Circuit is enabled
Circuit is enabled	Data is available or on incoming call
Data is available or on incoming call	Circuit is enabled
Data is available or on incoming call	Data is available or on incoming call

The valid combinations of connection initiation parameters for circuit group members and circuit group backup members using V.25 bis adapters are shown in the shaded fields of table 4. Pool circuit group members can *only* be configured to connect when “data is available or on incoming call”.

As in the case of manual adapters, when the connection initiation paramenter is set to “circuit is enabled” on both sides, the circuit will never hang up. The user will incur charges for a permanently open circuit!

Connection Controls—V.25 bis Adapter

V.25 bis circuits provide the same connection controls as described previously in the section “Connection Controls—Manual Adapter”. In addition, call restrictions may be configured with V.25 bis circuits. This involves specifying a list of originating phone numbers for which incoming calls will be accepted. Note that these call restrictions only work on ISDN networks with automatic number identification (ANI).

Sample Networks

The following sections show some of the applications for switched circuits and the methods for using circuit-switching networks on HP routers.

Low-Volume, Infrequent Transmissions

Figure 11 shows a network in which the traffic volumes are low and transmissions are infrequent. End-of-business-day transmissions to update databases is an application that may have only modest data transmission requirements. In this scenario, several (or many) sites are networked to a central site using switched services.

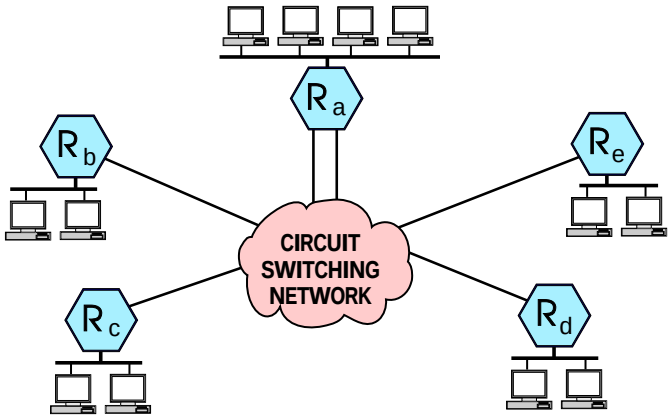

Figure 11. Routing Using Switched Network Services

The switched circuits at the remote sites are configured as primary circuits. The switched circuits at the central site are pool circuits. Note that there are more remote circuits than there are central-site circuits. The balance between remote-site dial-access circuits and central-site access circuits is determined by the frequency of access and by users' tolerance for waiting for an available access circuit. Switched circuits at the central site may be configured (by the Telco/PTT) within a hunt group (also known as a rotary), in which the remote sites dial a single number and obtain the first available circuit.

Alternatively, the central-site switched circuits may be configured with individual phone numbers. Remote sites (under V.25 bis control) would then have several possible numbers to dial to get around busy conditions.

Switched-circuit-only networks can be a very cost-effective alternative when used for several hours or less each day. To ensure that the switched circuit is open (connected) for the minimum possible amount of time, it is essential to disable the dynamic routing protocols (RIP and OSPF) and configure *static routes*. Dynamic routing protocols generate periodic traffic that will hold switched circuits open permanently!

Note Pool circuits may be used only with IP protocol.

NxN Backup

NxN backup refers to a configuration in which every leased-line circuit has its own dedicated backup circuit, for example, in a central site. Dedicated backup circuits are provided using the *backup* circuit type. Either manual adapters or V.25 bis adapters can be used to provide the backup circuits. Consider the network shown in figure 12. Each leased line has a dedicated backup. When link 1 fails, R_b dials R_{a1} (or vice versa). When link 2 fails, R_c dials R_{a2}. A backup circuit is required for each leased line.

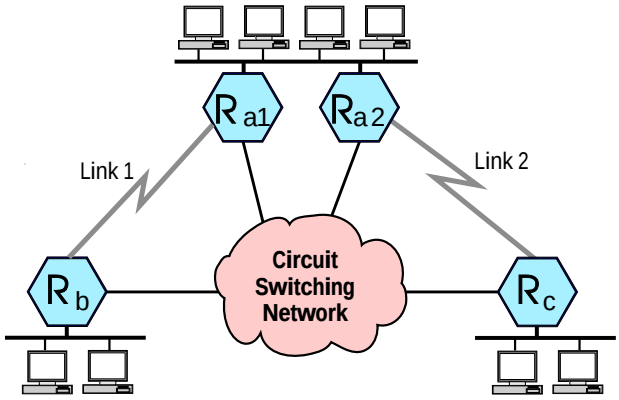

Figure 12. NxN Backup for Leased-Line Network

Dedicated backup circuits are easy to set up and may be used with all the routable protocols. There is no need to configure static routes, since the circuit will be disabled once the primary circuit (leased line) is restored. The disadvantage is that a separate switched circuit, a router port, and an adapter or DSU/CSU must be provided for each leased line. Compare this backup solution to 1xN backup in the next section.

1xN Backup—Star Network (IP Only)

Figure 13 shows a router network that uses leased lines as the primary circuits and dial up lines for backup purposes. (The backup circuits are ***not*** configured as backup members of a circuit group.) In this example, two switched circuits at the central site are used to back up four leased-line circuits (1 for 2 backup). R_{a1} and R_{a2} are attached to the central site LAN and connected to routers at four remote sites. The backup circuits become active only in the event of a primary link failure. Additionally, the backup circuits are only used to connect from a remote site to the central site (no remote-site-to-remote-site connectivity). In this example, the backup circuits only work for IP protocol, since only IP has the capability to map a phone number to a protocol address.

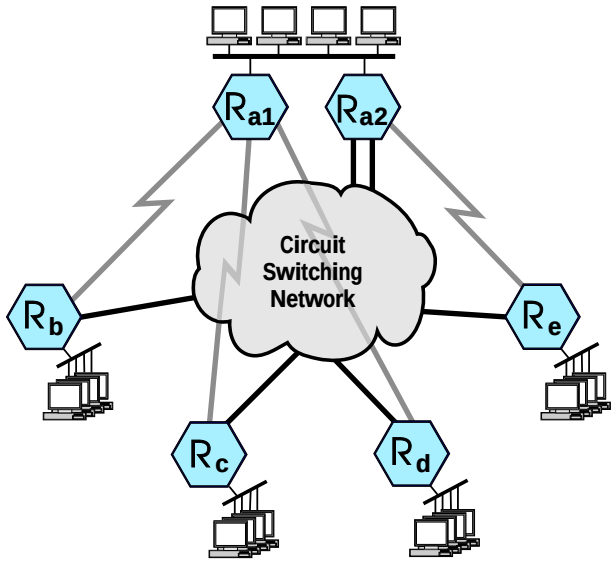

Figure 13. Leased-Line Network with Circuit-Switching Backup

The network in figure 13 is set up as follows:

- The leased lines and Ethernet LANs are configured normally.
- The switched circuits at the remote sites can use manual or V.25 bis adapters that can be configured to connect when the “circuit is enabled” or when “data is available”. The use of V.25 bis adapters configured to connect when “data is available or on incoming call” is preferred since it provides the most flexibility and control. The switched circuits are configured as primary circuits (circuit group members).
- The switched circuits at the central site must use V.25 bis adapters and must be configured to connect when “data is available or on incoming call”. These circuits must be configured as pool circuits (pool circuit group members). Additionally, IP pool circuit map entries must be configured (to map IP addresses to phone numbers).
- Switched circuits must be configured using static routes. Each leased line is favored over its switched counterpart by ensuring that leased-line circuit costs are set to result in the lowest-cost path.
- All switched circuits are members of the same IP network or subnetwork. Each leased line is a separate network or subnetwork.

1xN Backup—Meshed Network (IP Only)

Figure 14 shows another case of 1xN backup. This time the network is a packet-switching network (frame relay or X.25 point-to point). Unlike the network in the previous example, this network does not have a central site. Clients and servers exist at each of the sites. Thus, traffic flows between each of the sites. The network backup objective is to open a single switched circuit when one of the router's packet-switching network-access circuits fails and still to maintain the ability to communicate between all of the sites.

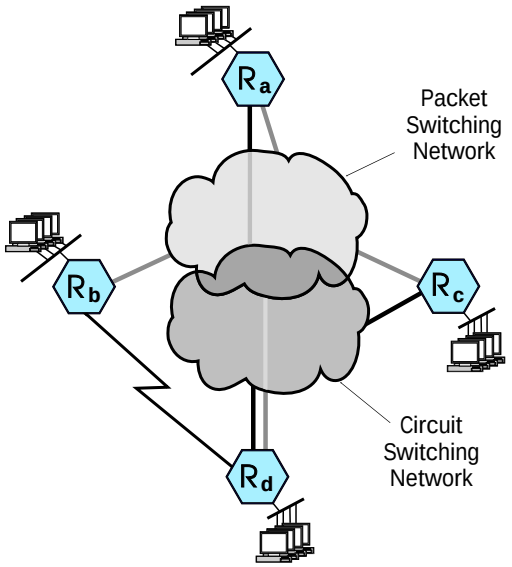

Figure 14. Packet-Switching Network with Circuit-Switching Backup

The network in figure 14 is set up as follows:

- Connections to the packet-switching network are configured in a single subnet using either RIP or OSPF. OSPF is preferred since it detects failures and switches around failures faster than RIP.
- Connections to the circuit-switching network are configured in a single subnet (a different subnet than the one used for the packet-switching network).
- Router D (R_d) is the central site from the perspective of the dial-backup scheme. R_a and R_c have static routes to each of the other Ethernets; R_d is the next hop for all of the static routes to the remote Ethernet networks. The point-to-point circuit between R_b and R_d has been installed to provide a backup path in the envent of a failure in R_d's packet-switching access circuit. A new feature in the next HP router software release—conditional static routes—allows you to define a new type of static routeand removes the need for the point-to-point circuit.

- The static routes defined for the circuit-switching network must have higher assigned costs than the routes used on the packet-switching network but equal preference. Thus, packet-switched routes are selected over circuit-switched routes to maximize throughput and minimize cost.
- Detailed configurations for this network example are available from HP on request.

Bandwidth on Demand

HP routers support V.25 bis extensions defined by Cisco Systems, Inc., and Ascend Communications, Inc. The extensions allow control of several parameters used for establishing communications on ISDN networks. The inverse multiplexer parameters that can be controlled by HP routers include:

- Per-channel bandwidth (56 Kbit/s, 64 Kbit/s, etc.)
- Minimum channels to aggregate (the minimum number of ISDN channels to open when data is available).
- Maximum channels to aggregate (the maximum number of channels to open regardless of the amount of data to transmit).

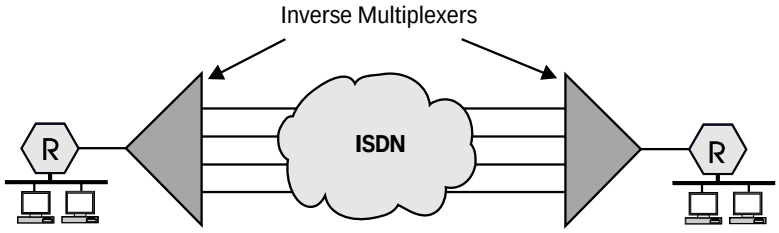

Figure 15. Inverse Multiplexers Used for Bandwidth on Demand

Figure 15 shows two routers that are connected to inverse multiplexers that are attached to an ISDN network. The inverse multiplexers open circuits as needed. When the traffic volume from the routers increases, additional circuits are opened. As the traffic volume decreases, circuits are closed.

Inverse multiplexers are available from several companies including Ascend, Simplex, and Presticom. Additional features can include data compression, Switched 56 and PSTN switched-circuit types, and leased-line support with incremental switched-circuit bandwidth.

Application Recovery

One question that often arises when considering router networks with backup mechanisms is: “Will my application recover transparently after a primary circuit fails?”. There are two requirements that must be met for network applications to recover when failures in router-based networks occur:

- First, the router network must have another route to the target system or network, and the routers must be able to detect failures and establish alternate routes quickly. The routing protocols, as well as the Spanning Tree Protocol, meet the fast switching criteria.
- Second, the network architecture on which the application runs must have recovery mechanisms (usually in the transport layer) to tolerate brief communication failures.

Applications running at the time of the failure may or may not recover transparently. IP applications will almost always recover transparently. NetWare users may receive an “Abort, Retry, Fail” message. (Answering “Retry” will generally succeed in restoring communications.) DECnet applications may fail when default network parameters are used. However, you may modify DECnet transmission parameters to increase the applications’ tolerance of temporary network failures.

Case Study

ISDN Wide Area Network Design: Dry Creek Joint Elem. School District

Larry Angus, Network Consultant
Hewlett-Packard Company

Organization Overview

The Dry Creek Joint Elementary School District is recognized as the fastest growing school district in the state of California. To maintain continuity of service with student increases averaging 47% for the last two years, the district has embarked on an aggressive use of technology. This has included the design and implementation of local area networks for each of its three school sites, all connected to the district office through a wide area network.

Business Need

Historically, the Dry Creek Joint Elementary School District has been a small single-school district. It was established in 1876 and maintained a stable population of between 100 and 200 students for most of its first 100 years of existence. Around 1988, significant residential development occurred within the district, and the student population began to increase by 400 to 600 per year. The current enrollment is almost 2,300 students.

With this rapid increase in student enrollment, the district constructed two new state-of-the-art elementary schools, which include the necessary information infrastructure to provide for computers in the classroom, integrated computer labs, integration of computers in libraries, cable television in each instructional area, telephones in each of the classrooms, and office spaces for teachers. Dry Creek has implemented year-round educational programs in all three schools.

To provide for the exchange of administrative and instructional information, the three school locations require connection with the district office. In partnership with Hewlett-Packard, the district has selected a two-phased strategy using extended local area network (LAN) technology. Phase I implements access to student administrative applications, electronic mail, and other office automation applications for a core of trained administrators and clerical staff.

During the 1993-94 school year, the district will begin phase II by designing and implementing an integrated instructional computer network, tying all school instructional areas together in a wide area network with access to external resources such as the Internet.

Applications

Dry Creek's enterprise network must support office automation applications such as electronic mail, word processing, graphics, and electronic spread-sheets, as well as host- and server-based business and student administration applications. The wide area network also must handle the additional load of instructional traffic, such as electronic mail and Internet access, when those services are brought online.

Network Topology

The current network is a star topology with the district office as the center hub. Currently, the entire district is serviced by Roseville Telephone Company from one central office. Roseville Telephone Company can provide either 56 Kbit/s leased digital circuits or basic-rate Integrated Services Digital Network (ISDN) services. The district's criteria for its wide-area service are that it must be cost-effective, provide high performance, be supportable, and provide maximum uptime.

ISDN and point-to-point circuits are the only carrier options available to the district. The service costs of both options are shown in table 1.

Table 1. WAN Cost Comparison (Q1/1993)

Circuit Type	Monthly Cost	Installation Cost
56 Kbit/s point-to-point (4 locations)	$324.30	$ 3,720.00
ISDN (4 locations)	110.50	849.75

The traditional point-to-point extended-LAN topology that was considered is shown in figure 1.

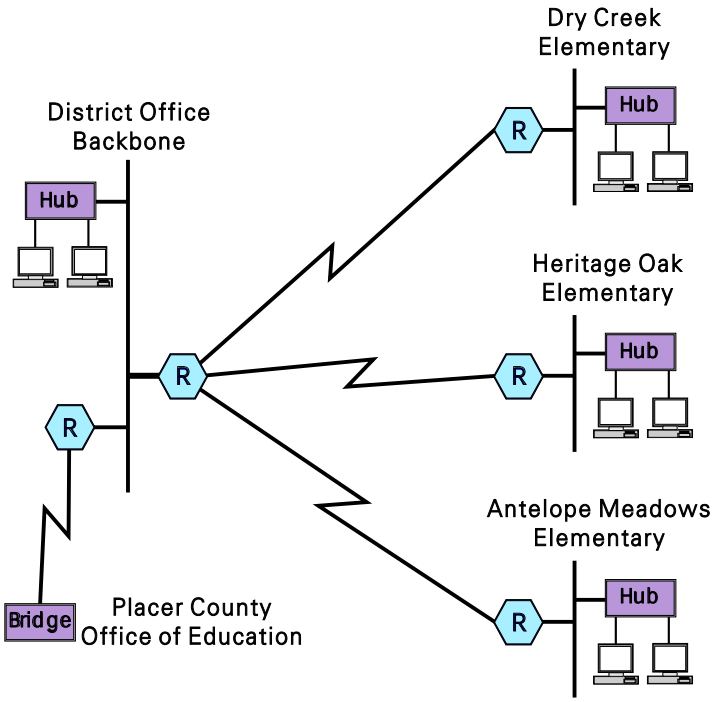

Figure 1. Point-to-Point Network Design Option

ISDN

Traditionally, the point-to-point network depicted in figure 1 would be implemented with leased 56-Kbit/s circuits; however, the low cost of ISDN required a thorough investigation to see whether this option was appropriate for Dry Creek Schools.

ISDN is a switched technology provided as a basic-rate service or a primary-rate service. The basic-rate service, called 2B + D, provides two 64-Kbit/s "B" ("bearer") channels. The primary-rate service, in North America called 23B + D (or in Europe 31B + D), provides 23 channels for a variety of voice, video, and data connection options. Each service includes one control or messaging "D" channel (16-Kbit/s for basic service). Basic-rate service (2B + D) is all that Dry Creek School District needs. The implementation considerations for basic-rate ISDN only are discussed in this article.

When comparing cost and speed to digital 56-Kbit/s services, ISDN seems to be the better choice, but ISDN has potential problems. The terminal adapter equipment is currently more costly than DSU/CSU equipment, and ISDN services are not always available in all areas of a particular local exchange carrier or between carriers. In many cases, a digital link may be available between two end points, but the call may start out as ISDN with out-of-band signaling, be converted to Switched 56 with in-band signaling, and be converted back to ISDN. In this example, there would be a loss of throughput because of the Switched 56 in-band signaling. Signaling System 7 is a standard network service that will eliminate this problem when it is universally implemented by all carriers.

Another design consideration is that there may be measured charges incurred for an ISDN data call, depending on the geography of the sites. Usage charges are regulated and will be applied to carrier network services, depending on how the customer is set up as a business group with a particular carrier. Extended LAN connections regardless of protocol tend to establish their connections for long periods. A good rule of thumb is that if there are toll charges or measured-business-unit charges for voice connections between two end points, there probably will be such charges for ISDN data. Obviously, all ISDN charges must be considered when comparing the costs of ISDN to other digital services.

All of Dry Creek Joint Elementary School District is serviced as one business group by Roseville Telephone Company. With no usage charges, ISDN meets all of the district's design criteria. The network as implemented is shown in figure 2.

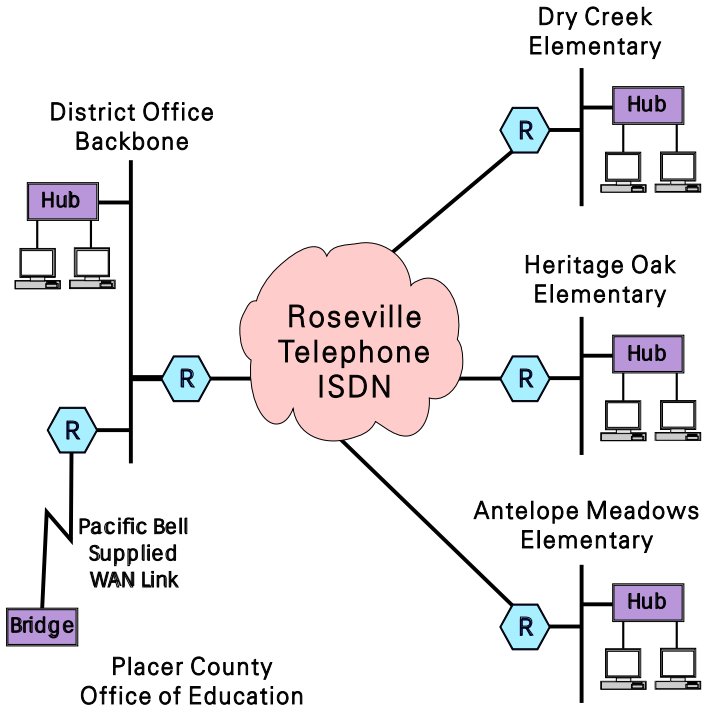

Figure 2. ISDN Network Design Option

HP Routers

Initially, the district's network will be implemented as a bridged network bridging IP. HP routers were selected for use in the Dry Creek design for the following reasons.

1. HP routers can be supported on a standard contract with a guarantee of 4-hour response time for repair.

2. HP routers support ISDN as of release 5.74.

3. The HP Router SR represents a lower-cost design when compared to multiple bridges at the home site.

4. HP routers can support link speeds from 56 Kbit/s to T1, should higher bandwidth be required in the future.

5. HP routers provide more complete troubleshooting capabilities—NCL, statistics screens, and HP OpenView Network Management—than do bridges.

With HP router operating system release 5.74 and later, ISDN can be configured for manual or automatic operation. The automatic operation provides for very sophisticated call control using the CCITT V.25 bis call-connect protocol. The Dry Creek implementation uses point-to-point connections that are semi-permanently established, requiring only simple manual operation to be configured on the routers.

ISDN Configurations

ISDN is a switched/dialed technology and requires that one end point of a point-to-point connection be configured to originate and that the other end point be configured to answer. Each end point is assigned an ISDN telephone number by Roseville Telephone Company. The configurations for the routers and ISDN terminal adapters are shown here.

District Office Router

Basic configuration: Host-only bridging enabled, TFTP enabled, SNMP enabled, and Telnet enabled. The district office router is the originator and causes the terminal adapter to make the call when the router is booted and its DTR goes high.

HP Router SR ISDN Circuit Configuration:

```
Circuit Name:  WAN2
Auto Enable:  YES
Quality of Service:  LLC1 (datagram)
Circuit Type:  ISDN Manual Adapter
Min Frame Spacing:  2
Connect Retry Count:  3
Connect Wait Time (sec):  15
Delay After Connect Failure (min.):  .5
```

District Office Terminal Adapter (TA)

AT&T 7500B Data Module operating in synchronous DCE mode acting as the originator.

Configuration:

```
Set Autodial = YES
Set busy out = OFF
Data Rest = OFF
Set DTR = FOLLOW
Set Duplex = FULL
Set mode = CS
Set speed = 64000
```

The ISDN telephone number for the remote site is configured in the terminal adapter's call table; the autodial feature engages when the router's DTR goes high.

Remote Router

Basic configuration: Host-only bridging enabled, TFTP enabled, SNMP enabled, and Telenet enabled. The remote-site router is the answerer and causes the terminal adapter to wait to answer the call when the router is booted and its DTR goes high.

HP Router FR ISDN Circuit Configuration:

```
Circuit Name:  WAN1
Auto Enable:  Yes
Quality of Service:  LLC 1 (datagram)
Circuit Type:  ISDN Manual adapter
Min Frame Spacing:  2
Connect when:  Circuit is enabled
Connect retry count:  3
Connect wait time (sec.):  0 [1]
Delay after connect failure (min.):  .5
```

[1] It was discovered during testing that this parameter should be set to zero to ensure that the "answer" end point is ready to reconnect immediately in the event of a power loss or component failure.

Remote Office Terminal Adapter (TA)

AT&T 7500B Data Module operating in synchronous DCE mode acting as the answerer.

Configuration:

```
Set Autodial = NO
Set busy out = OFF
Data Rest = OFF
Set DTR = FOLLOW
Set Duplex = FULL
Set Mode = CS
Set Speed = 64000
```

Performance

ISDN provides clear 64-Kbit/s channels for each point-to-point connection. This bandwidth will meet Dry Creek's needs for approximately two to three years. Basic-rate terminal adapters that provide inverse multiplexing are becoming available. The inverse multiplexing feature will allow both B channels to be combined for 128 Kbit/s of bandwidth. This provides a growth path that will be more than adequate for Dry Creek's needs for at least three years.

Shining a Light on FDDI

FDDI—the 100-Mbit/s Fiber Distributed Data Interface networking technology—is the solution for many of the new problems presented by changing corporate networks:

- Groups that previously had no need for communication now want network connections.
- Existing token ring and Ethernet backbones that are interconnected are now reaching their capacity.
- Applications require increasing bandwidth, security, and fault tolerance.

These are examples of some of the network requirements that are met by FDDI. At 100 Mbit/s, FDDI allows high-speed interconnection of all the LANs in an organization's network. FDDI easily overcomes the performance limitations of 10-Mbit/s Ethernet and 16-Mbit/s token ring networks. In addition, FDDI provides the security of fiber-optic cabling as well as the fault tolerance that is built into the network design. As a result, FDDI is becoming widely accepted, especially as a network backbone technology.

Technology Overview

Fiber Distributed Data Interface (FDDI) is an ANSI and ISO specification (X3T9) for the transmission of data at high speeds, typically 100 Mbit/s, using fiber-optic cable as the transmission medium. Optical fiber technology offers networks a great degree of flexibility in bandwidth and topology design. Fiber-optic cable also offers excellent noise immunity, and is virtually impossible to tap.

FDDI is a token-passing technology that uses a timed-token protocol to guarantee network access between network stations (network devices and end nodes). Figure 1 shows a standard dual-attached ring. Network access is negotiated between stations at initialization and at every time a new node is added to the network.

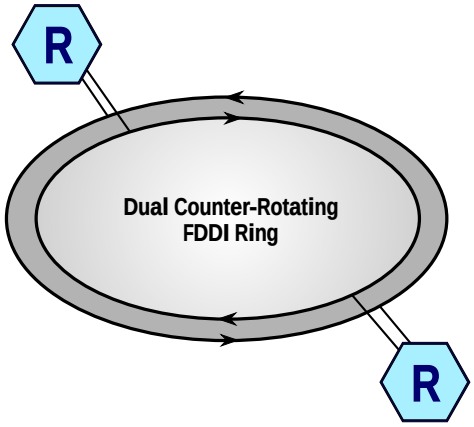

Figure 1. FDDI Ring

The network backbone is constructed of stations interconnected by two counter-rotating rings. These rings are two pairs of fiber-optic cable to which each device is attached. Cable lengths between stations can be anywhere from 2 kilometers (km) with multimode fiber, to 60 km with single-mode fiber. The total ring length cannot exceed a maximum of 200 km. During normal operation, the first ring is the primary data carrier, and the second acts as the backup. This offers the network a greater degree of redundancy and fault tolerance.

FDDI and the OSI Model

The FDDI standard is made up of four distinct parts:

- Physical Layer Medium Dependent (PMD) and
  Single Mode Fiber Physical Layer Medium Dependent (SMF-PMD)
- Physical Layer Protocol (PHY)
- Media Access Control (MAC)
- Station Management (SMT).

Although not part of the FDDI standard, the Logical Link Control (LLC) is required by FDDI to assure transmission of user data. These standards define the 100-Mbps fiber-optic dual counter-rotating FDDI ring.

The PMD, SMF-PMD, and PHY are equivalent to the physical layer of the OSI model (see figure 2). The PMD and SMF-PMD correspond to the lower portion of the physical layer. The PMD defines the media requirements for multimode fiber, such as fiber-optic cable, connectors, and driver receiver operation for FDDI stations. (Cables, connectors, and receivers are discussed below in the “Cabling” section.) The SMF-PMD defines similar requirements for single-mode fiber-optic media. The PHY corresponds to the upper portion of the physical layer and defines the symbol set, link states, encoding/decoding, clocking, and framing.

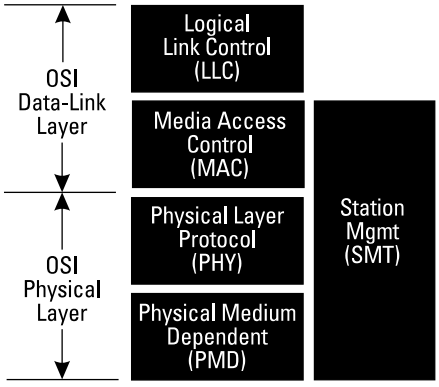

Figure 2. FDDI Protocol Stack

The MAC corresponds to part of the data-link layer of the OSI model. The MAC standard defines the token-passing method as the means for acquiring access to the ring. It is responsible for frame and token construction, sending and receiving frames on the FDDI ring, and delivering LLC frames.

The LLC service provides the transmission of a frame of data between two stations. LLC frames carry user information to stations on the ring and also to the extended LAN.

FDDI also has a process that defines protocols for managing the PMD and SMF-PMD, PHY, and MAC, called Station Management (SMT). SMT defines facilities for connection management, station configuration, error recovery, and the encoding of SMT frames. MAC and SMT frames carry data and control information for the operation and management of the FDDI network.

The MAC, PHY, and PMD standards were approved by ANSI and ISO by 1990. The SMT standard was approved in 1991. FDDI is the only LAN with extensive management capabilities as defined in the SMT.

How FDDI Works

FDDI is based on two counter-rotating 100-Mbps fiber-optic token-passing rings. If one ring should fail, FDDI automatically becomes a single, not dual, FDDI ring. The rings consist of point-to-point connections between adjacent stations. Stations negotiate for ring access at initialization and when new stations are added to the ring. The two rings act as a primary carrier and a backup (secondary) carrier. Data flows in opposite directions on the two rings.

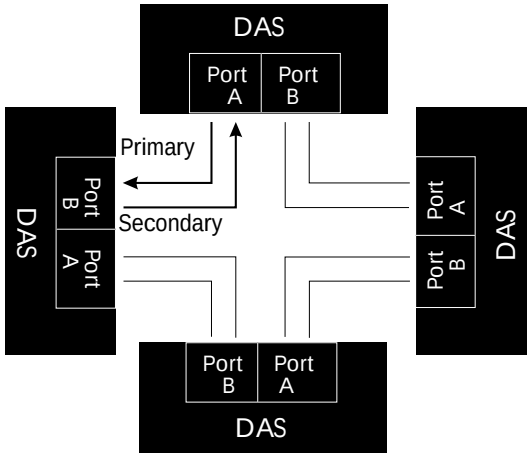

Figure 3. Dual Attachment Stations on an FDDI Network

FDDI Stations

There are three types of FDDI stations: Dual Attachment Stations (DAS), Single Attachment Stations (SAS), and concentrators. See figure 3.

Dual Attachment Stations are devices that connect to both the primary ring and the secondary ring simultaneously. Single Attachment Stations are devices that connect to the ring through the primary ring only, and are usually attached to the ring through concentrators.

Dual Attachment Stations have two ports, one labeled port A, the other port B. On each port, there is a primary and a secondary ring connection; port A is primary IN secondary OUT, port B is primary OUT secondary IN. If either port fails, the other port then becomes primary IN and OUT on the same port.

Concentrators connect directly to the backbone of the ring, and provide indirect access to the ring for other devices. Concentrators support multiple DAS and SAS connections, and add a degree of fault tolerance to the network by isolating end nodes from the ring.

Ring Access

When a station is first attached to the FDDI ring, it undergoes an initialization process. During this initialization process, a newly inserted station initiates a connection management process (CMT) with its neighbor. CMT tests port types, performs physical port tests or link confidence tests, and transmits link-state characters upon completion of the tests. Stations initiate CMTs between each upstream and downstream neighbor.

Upon successful completion of the initialization phase, the ring begins the claim process to generate a token. The claim process begins with all stations negotiating a target token rotational timer (known as TTRT). Each station makes a bid for the timer, and the one with the lowest value wins the right to initialize the ring by inserting the token. If two stations happen to make the same bid, the station with the highest address wins the right to generate the token. Once the claim process has finished, the ring enters a steady state, to be disrupted only by the insertion of another station, or by a break in the ring.

Access to the FDDI network is controlled by a token that circulates the primary ring. When a station has data to transmit, it must capture the token before it can transmit data onto the ring. Only one token may exist on the ring at any one time. If the token is lost, the ring will enter a phase known as beaconing (a state in which the beaconing station controls all access to the ring). If more than one token has been generated, the ring is then scrubbed (scrubbing is used to remove all data from the ring).

Ring Wrap

Ring wrap is one of the three techniques that FDDI uses in order to ensure fault tolerance. If any station on the dual ring fails or if the cable is damaged, the dual ring automatically wraps into a single ring as shown in figure 4.

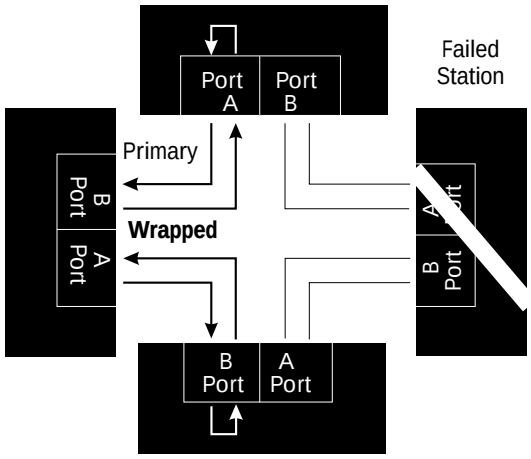

Figure 4. Ring Wrap

Multiple failures within a single dual-ring network present other challenges. The likelihood of multiple failures greatly increases as the number of nodes within an FDDI network increases. If a single failure occurs, the ring wraps and becomes a single ring. If another failure should occur, the ring is segmented and separated into two rings.

Optical Bypass

Optical bypass can be used for fault tolerance to prevent ring segmentation. Optical bypass switches maintain connectivity of the FDDI ring in the absence of power or during fault conditions in a station. Stations bypassed by optical bypass switches are effectively removed from the ring. Figure 5 shows an inoperative station switched out by an optical bypass switch.

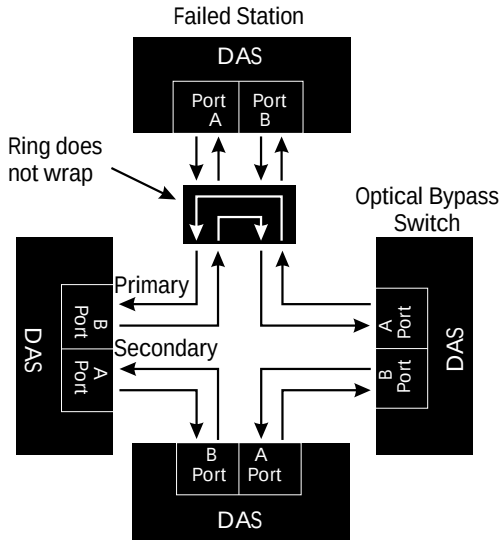

Figure 5. DAS with Optical Bypass Switch

Optical bypass switches are recommended when designing a multistation FDDI backbone. In a fault condition, a DAS will cause the ring to wrap over the secondary ring, isolating the bad section of the ring. If this happens to more than any single section of the FDDI ring, the network is sectioned into two independent rings. Stations that have an optical bypass device are physically switched out when a fault condition occurs. They do not prevent the network from segmenting in a cable fault condition.

There are several limitations of optical bypass that a network designer must be aware of:

- When the bypass switch bypasses a station, the station is effectively removed, possibly allowing the maximum segment length to exceed 2 km.
- The integrity of the ring is only as good as the mechanical integrity of the bypass switch.
- Optical bypass switches can produce up to 2.5 dB optical power loss. This loss must be considered when calculating attenuation. (Attenuation and optical budget are discussed below.)

Dual Homing

Dual homing is a fault tolerance technique used for redundancy in concentrator environments where continuous uptime is crucial. Dual homing allows a DAS to use one link as a backup link for redundancy purposes. Thus, one of the two attachments is active at any given time. Figure 6 shows a dual-homed device attached to concentrators.

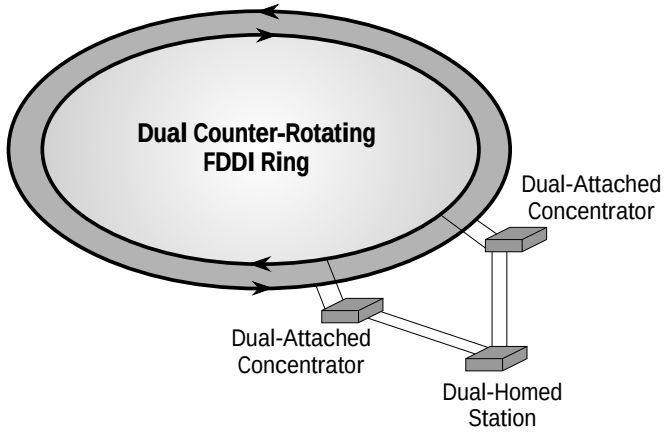

Figure 6. Dual-Homed Station

A station can be dual-homed to the same concentrator or to two different concentrators. When it is homed to two different concentrators and the primary concentrator fails, the DAS automatically enables the backup link to the secondary concentrator.

Cabling

Optical fiber has many advantages over traditional copper cabling. Fiber doesn't emit electrical signals and is immune to electrical interference, making it secure and reliable. It is easy to manipulate because it is light-weight. Important specifications of fiber-optic cable are attenuation and the optical signal wavelength the cable can carry.

Two types of transmission media are currently defined by the FDDI standards: multimode fiber and single-mode fiber. Both operate at 100 Mbit/s. Multimode means that multiple rays of light can enter the fiber from different angles. Multimode uses light-emitting diodes (LEDs) that convert electrical signals into light and transmit the light into the fiber-optic cable. Single mode means that only one ray of light is allowed to enter the fiber. Single mode uses laser diodes (LDs) to convert electrical signals into light and transmit the light into the fiber-optic cable. The HP routers do not support single-mode fiber. They do support multimode fiber—only at 1300 nanoseconds.

Attenuation and Optical Budget

Signal attenuation through optical fiber is important with FDDI. It describes the amount of energy (optical power) that is lost as the light signal travels from the transmitter through the cable to the receiver. The longer the cable, the higher the loss of optical power. Energy loss is denoted in decibels (dB), an expression used to mathematically compare the power of two signals.

Attenuation is calculated by knowing the unit attenuation and the length of the link. The maximum cable attenuation is the power loss in the cable as well as any loss incurred by splices, connectors, and anything else connected to the cable. PMD specifies an optical power budget between any two stations of 11 dB. SMF-PMD allows for a range of power budgets that extends from a minimum of 10 dB to a maximum of 32 dB.

To calculate the cable loss between two adjoining stations, use the formula shown below:

Attenuation = (Cable Len (km) x Attenuation/km) +
(Splices x Attenuation per Splice) +
(Connectors x Connector Attenuation) +
(Loss at Transmit MIC)

For example:

2 km fiber x 2.5 dB/km	5.0 dB
2 splices x 0.25 dB/splice	0.5 dB
1 connector x 0.5 dB/connector	0.5 dB
Loss at transmit MIC	+0.5 dB
Total link attenuation	6.5 dB

Once you know the attenuation, you can calculate the remaining power available for the network. The maximum allowed loss is 11 dB. The formula is shown below:

Power Available = Optical Power Budget – Attenuation

For example:

Optical power budget	11 dB
Attenuation	– 6.5 dB
Total available	4.5 dB

After calculating the dB loss of a particular link, it is important to compare the resulting power with receiver sensitivity. If the receiver is not sensitive enough, poor link quality will result.

Optical Transmitters and Receivers

Transmitters convert data from electrical signals to light. The receiver converts the light signals back to electrical signals. Receivers contain photo detectors that convert incoming optical signals back into electrical signals.

Optical transmitters convert modulated electrical signals into modulated light signals that are transmitted through the fiber-optic cable. In multimode fiber cable, transmitters are light-emitting diodes (LEDs) that convert electrical signals into light signals. In single-mode fiber cable, transmitters are laser-quality light-emitting diodes or laser diodes (LDs).

Internetworking with HP Routers

HP's FDDI link interfaces are compliant with the following ANSI X3T9.5 standards:

- Physical Medium Dependent (PMD)
- Physical Layer Protocol (PHY)
- Media Access Control (MAC)
- System Management (SMT)

They also comply with RFC 1188, which specifies transmission of IP datagrams over FDDI media, and with IEEE 802.1 Parts D & H.

Physically, they provide a single MAC connection and support both Class A and Class B attachments to FDDI media. A Class A attachment requires two physical connectors and provides connectivity between a Dual Attachment Station (DAS) and the FDDI primary and secondary rings. A Class B attachment uses a single physical connector and provides connectivity between a Single Attachment Station (SAS) and the FDDI primary ring or an FDDI concentrator. HP's FDDI link interfaces only support multimode 50-micron or 62.5-micron graded-index fiber-optic cable.

FDDI Routing

Since routers are media-independent, they encapsulate packets to conform to the media type in use. RFC 1188 specifies 802.2 LLC with SNAP (subnetwork access protocol) when routing TCP/IP traffic over FDDI media. RFC-1188-conforming routers provide IP routing between source and destination end systems on Ethernet or 802.3 LANs across an FDDI backbone, as well as to destination systems directly connected to the backbone.

Other routable protocols are translated on the FDDI ring according to rules defined by RFC 1188 and IEEE 802.1 bridge standards. Adherence to such rules allow interoperability between multiprotocol routers and translating bridges that may be attached to the same FDDI ring. The protocols supported on FDDI include:

- IP
- Novell IPX
- Appletalk Phase 2
- DECnet Phase IV
- Xerox XNS
- HP Probe

Encapsulation Bridging

There are many protocols that are non-routable and therefore must be bridged. These non-routable protocols must be encapsulated in an FDDI frame in order to be transported across the ring. Encapsulation is the bridge's implementation that enables interconnection of similar networks over an FDDI network. Figure 7 shows how non-routable protocols might be encapsulated across an FDDI network. Within such a topology, a bridge encapsulates the original Ethernet/IEEE 802.3 frame into a new message type (in this case, an FDDI-specific packet) for travel across the FDDI ring. At the destination, the message is removed from the FDDI packet and transmitted in its original form.

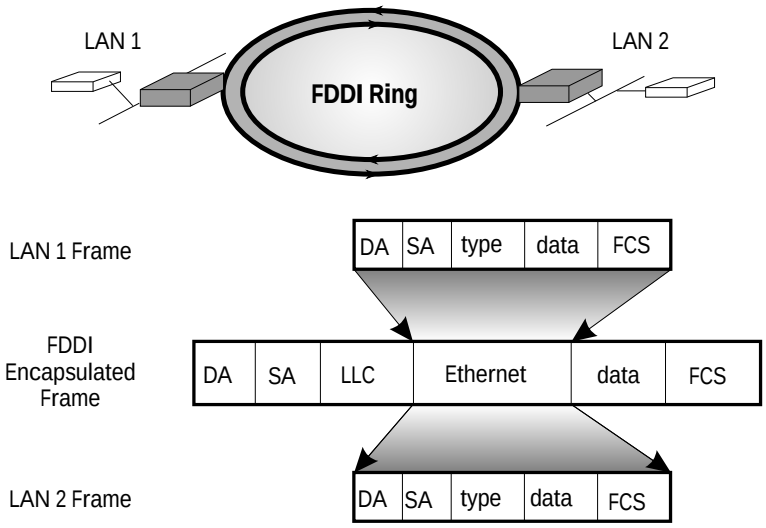

Figure 7. Encapsulation Bridging

It is important to note that encapsulation is not the preferred method of transporting data across an FDDI network. In addition, no standard method of encapsulation exists, preventing multivendor interoperability.

FDDI Translation Bridging

To ensure multivendor interoperability, the bridge protocol should be based on the IEEE 802.1 Spanning Tree Protocol. Translation is required when bridging between LANs with different data-link-layer characteristics. For example, forwarding from an Ethernet to an 802.3 end station requires a translation. When bridging Ethernet or 802.3 packets to FDDI media, a MAC-layer translation is also required.

The 802.1 standards define how a received packet is formatted for transfer across an intervening LAN (802.x or FDDI) and presented to a destination LAN. In general, Ethernet Version 2 frames are encapsulated according to rules specified in RFC 1042, and then converted back to Ethernet format. IEEE 802.3/802.2 LLC and 802.3/LLC+SNAP remain unchanged when transferred across an intervening LAN. Table 1 summarizes translation rules. Table 2 summarizes protocol-specific translation.

Table 1. Translation Rules for Encapsulation

Source Network Encapsulation	Translation Rules	Destination Network Encapsulation
Ethernet	RFC 1042 (SNAP OUI 00-00-00)	Ethernet
802.3/LLC + SNAP	FDDI LLC + SNAP	802.3/LLC + SNAP
802.3/802.2 LLC	FDDI 802.2 LLC	802.3/802.2 LLC

Table 2. Protocol Translation and Encapsulation

Protocol	Source Type	Translation Rules	Destination Type
XNS	Ethernet	RFC 1042	Ethernet
DECnet Phase IV	Ethernet	RFC 1042	Ethernet
Novell	Proprietary	RFC 1042	Ethernet
Novell	Ethernet	RFC 1042	Ethernet
AppleTalk Phase 2	LLC+SNAP	LLC + SNAP	LLC + SNAP
Source Routing	802.2 LLC	802.2 LLC	802.2 LLC
IP	Ethernet	RFC 1042	Ethernet
IP	LLC+SNAP	LLC + SNAP	LLC + SNAP

If an Ethernet frame generated on LAN 1 is destined for LAN 2, the frame is translated as shown in figure 8.

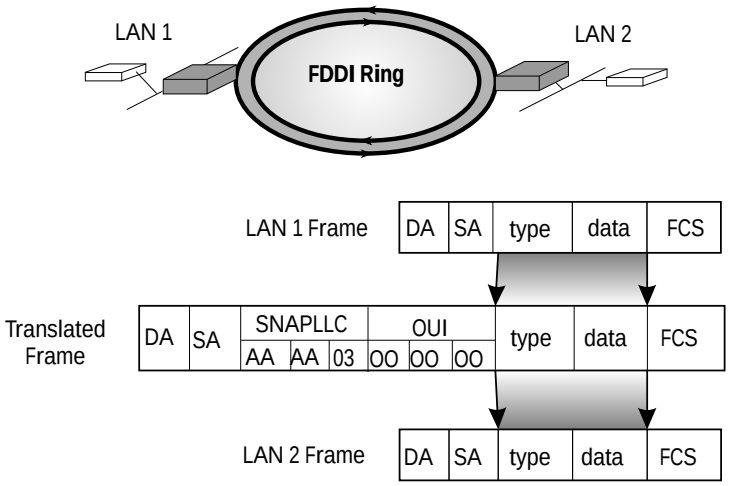

Figure 8. Translation Bridging

Translation by the bridge consists of:
1. extraction of addressing information from the Ethernet header
2. incorporation of address information into a newly generated FDDI MAC header
3. encapsulation of Ethernet data as specified in RFC 1088
4. FCS recalculation
5. addition of the FDDI MAC-level trailer

Standalone FDDI Rings

Network designers can take advantage of FDDI's high-speed capabilities without an FDDI backbone in a network topology. In the absence of a ring, a network designer can use the HP 27290A Router BR as a standalone FDDI network, interconnecting a maximum of two SAS devices or a single DAS.

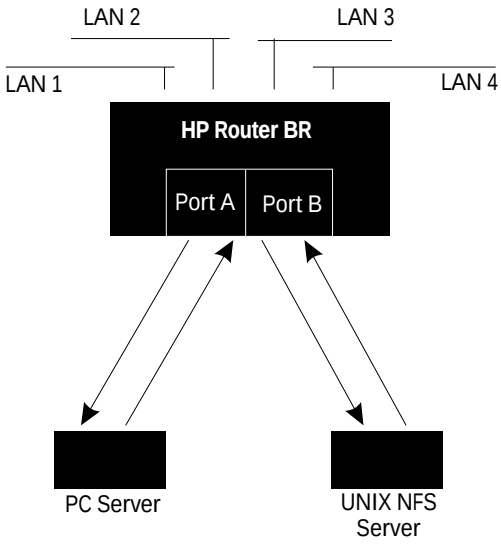

Figure 9. FDDI Super-Server Configuration

A file server is a very good candidate for a standalone FDDI network. In a standard configuration, a server would be connected directly to an Ethernet segment. Traffic between segments would be very heavy, with traffic on the server's segment being the heaviest. If servers were connected to the HP Router BR as shown in figure 9, two benefits would be:

- The traffic between segments would decrease draftically, effectively increasing the bandwidth of the Ethernet internetwork.
- Each server would have 100 Mbit/s bandwidth with which to serve all incoming segments.

In general, a device can be connected to the HP Router BR either as a DAS, creating a dual-attached ring, or as an SAS, creating a single-attached ring.

Using Synchronous Pass-Through to Consolidate Synchronous Traffic

Routers and high-speed WAN links have been used to create corporate internetworks. These internetworks connect LANs and provide communication among any systems and nodes attached to the LANs. There are many devices, however, such as SNA 3270 and 3770 terminals in most corporate networks, that are not connected to the extended LAN internetworks. Instead, these devices are typically connected to a parallel network. Parallel networks are expensive both for monthly connection charges and for network support.

Synchronous pass-through (also called sync pass-through) is a new feature of HP router software versions 5.70 and later that allows synchronous traffic (HDLC, SDLC, LAPB) to be consolidated with LAN-to-LAN traffic on HP-router-based internetworks. This application note describes how sync pass-through works and examines several sync pass-through applications.

Figure 1 shows an extended LAN internetwork where sync pass-through is used to connect synchronous interactive and printing terminals. Multiple IBM 3174 cluster controllers in a multi-drop configuration are connected to an IBM 3745 at site A over the extended LAN. Similarly, at site B, an HP 3000 running SNA IMF and SNA NRJE is connected to the IBM 3745 at site A using sync pass-through over the extended LAN. Using sync pass-through, a point-to-point circuit can be defined to convey synchronous traffic from one location to another in the extended LAN.

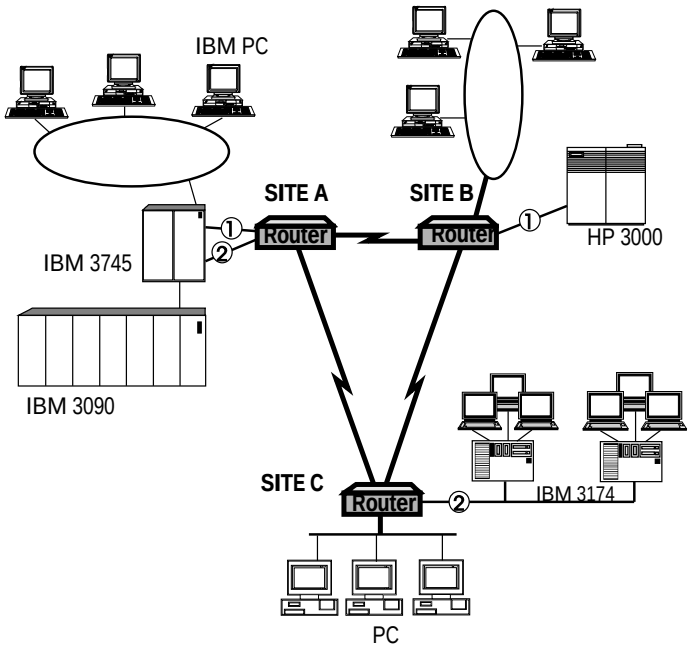

Figure 1. Synchronous Traffic Conveyed Through a High-Speed Extended LAN Network

How It Works

Synchronous ports on HP routers are used to provide the sync pass-through capability. Sync pass-through circuits are configured on the HP router using the point-to-point circuit type and the pass-through data-link-layer protocol. Once pass-through has been selected, the router software ignores the other circuit parameters. Additionally, the local and remote station addresses (MAC addresses) must be configured for sync pass-through to operate. Sync pass-through thus connects a device attached to one synchronous port on an HP router to another device attached to another router's synchronous port.

Sync pass-through traffic is conveyed through the extended LAN by the bridging service. Station addresses are assigned to sync pass-through ports by the user. In figure 1, station addresses are assigned to each of the synchronous circuits marked with a circled number. When configuring the site C sync pass-through interface, for example, the station address assigned to the local sync pass-through circuit (marked ②) is entered as the local LAN address. The station address assigned to the sync pass-through interface at site A (also marked ②) is entered as the remote LAN address.

Figure 2 shows another sync pass-through application. Sync pass-through is used to provide communication for a private X.25 network. Host A communicates with host B using the X.25 packet switches attached to the synchronous ports at site A and site B. The router ports labeled ① convey this switch-to-switch traffic.

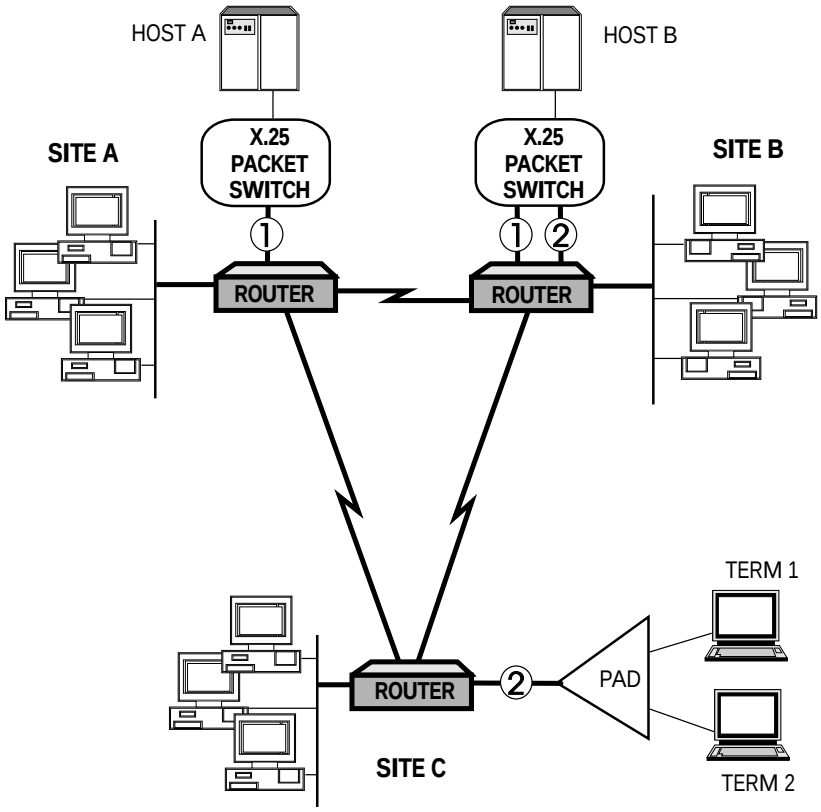

Figure 2. Bridging X.25 Traffic Using Sync Pass-Through

The PAD (packet assembler/dissembler) at site C is connected to the X.25 switch at site B using the sync pass-through ports labeled ②.

Traffic from terminal 1 to host B follows the path:

- terminal 1 to PAD
- PAD to site C router
- site C router to site B router
- site B router to X.25 packet switch
- X.25 packet switch to host B

Traffic from terminal 2 to host A follows the path:

- terminal 2 to PAD
- PAD to the site C router
- site C router to the site B router
- site B router to the attached X.25 packet switch
- X.25 packet switch back to site B router
- site B router to site A router
- site A router to attached X.25 packet switch
- X.25 packet switch to host A

Note that traffic from terminal 2 to host A does not use the router path from site C to site A. The PAD is connected to the X.25 switch at site B using sync pass-through.

Sync Pass-Through Encapsulation

The router constructs LAN packets by encapsulating the synchronous frames received on a sync pass-through port, as shown in figure 3. The local and remote station addresses assigned by the user are used as source and destination addresses. A special type field, recognized only by HP and Wellfleet routers, is added to mark the frame as a sync pass-through packet. The original CRC is removed.

These packets are then transmitted on the appropriate interface by the bridging service. Since the bridging service is used to convey sync pass-through traffic, bridging must be enabled on all router interfaces that may be required to convey sync pass-through traffic. Additional encapsulating protocol is added, depending on the type of circuit on which the sync pass-through packet is transmitted. (Figure 3 shows the encapsulation used on synchronous point-to-point circuits.)

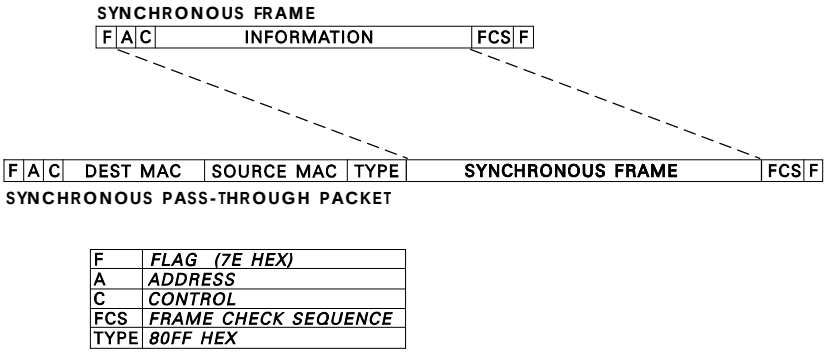

Figure 3. Sync Pass-Through Encapsulation on Backbone WAN Circuit

When the sync pass-through packet is received by the router that has the destination synchronous port, the encapsulating protocol is removed, a new CRC is computed, and the synchronous frame is transmitted to the destination synchronous device. The frame transmitted to the destination device is identical to the frame received by the first router.

Selecting Sync Pass-Through Station Addresses

To configure sync pass-through, a station address must be selected for each sync pass-through port. Care must be taken to avoid selecting a station address already in use on the extended LAN. Fortunately, the Ethernet/IEEE 802.3 standards provide a handy mechanism to accomplish this. Figure 4 shows the structure of a station address.

48-BIT ADDRESS FORMAT

I/G	U/L	46-BIT ADDRESS

I/G	0	INDIVIDUAL ADDRESS
I/G	1	GROUP ADDRESS
U/L	0	GLOBALLY ADMINISTERED ADDRESS
U/L	1	LOCALLY ADMINISTERED ADDRESS

Figure 4. Station Address Structure

Recall that station addresses are 48 bits in length. They are usually presented in a form such as XX-xx-xx-xx-xx-xx. Each “xx” indicates a single byte represented by two hexadecimal digits. “XX” indicates the most significant byte of the station address. The least significant bit of the most significant byte is the I/G bit. This specifies whether the address is an individual (I) or group (G) address. The I/G bit is the first address bit transmitted (although it is not the most significant bit), as shown in figure 4. The address selected for sync pass-through must be an individual address (I/G bit = 0).

The next bit (U/L bit) specifies whether the address is globally administered (U for universal) or locally (L) administered. This is the “handy mechanism” referred to above. Computer manufacturers always ship LAN networking products with globally administered (unique) addresses. The station address selected for use with sync pass-through should, therefore, be locally administered (U/L bit = 1). This should avoid any conflicts with any installed LAN equipment. The value of the most significant byte of the station address is, therefore, 02 hex (assuming I/G = 0, U/L = 1, all other bits = 0). Using the value 02 hex as the most-significant byte of the station address will avoid duplicating a station address used elsewhere in an extended LAN. 02-11-22-33-44-55 is an example of an individual, locally administered LAN address.

Synchronous Traffic Requirements

The type of synchronous traffic that may be conveyed using sync pass-through is limited to HDLC and HDLC derivatives such as SDLC and LAPB. The following three conditions must be met for sync pass-through to function.

- The synchronous port hardware recognizes synchronous traffic that uses flags (hex 7E character) as the idle line character. This means that flags are transmitted when the link hardware has nothing else to send. Flags are also used to delimit the beginning and end of synchronous frames.
- The maximum allowable sync pass-through frame size is 1600 bytes. Thus, the maximum allowable synchronous frame is 1600 bytes minus the encapsulating protocol size. (The size of the encapsulating protocol varies by link type but never exceeds 22 bytes.) Generally, a frame size greater than the maximum is not a problem, since most synchronous protocols do not use frame sizes greater than 1024 bytes.
- Additionally, the synchronous traffic must use the CRC-CCITT frame check sequence. This is given by the generator polynomial $x^{16} + x^{12} + x^5 + 1$

Synchronous traffic that does not conform to the above three requirements cannot be conveyed using sync pass-through. The most notable type of traffic that cannot be conveyed by sync pass-through is IBM BISYNC.

Physical Connections

Synchronous ports on HP routers physically function as DTEs.

Similarly, the synchronous ports on most devices, such as IBM 3270 cluster controllers and HP network interfaces, physically function as DTEs. Thus, the generally recommended practice is to connect these interfaces together using modems or modem eliminators (see figure 5).

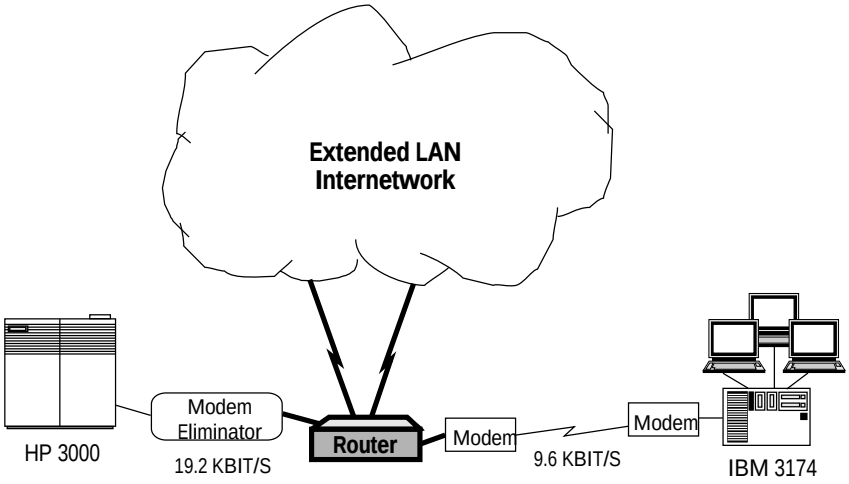

Figure 5. Modem Eliminators and Modems Used to Connect Synchronous Devices to an HP Router

A low-cost alternative to a modem or modem eliminator is a null-modem cable. A null-modem cable can be used instead of a pair of modems or a modem eliminator when the device to be attached to the router is located close by. Pin-outs for some null-modem WAN cables are shown in figures 6 and 7 for the HP routers. Figure 6 is for cables that can be attached directly to the router's 62-pin connector, and have either a V.35 or RS-232/V.24/V.28 interface. Making these cables requires a DB-62 connector on the router end. Figure 7 is for a V.35 cable that can be attached to the V.35 WAN cable that can be ordered with the HP router. Making the cable requires two standard M/34 V.35 connectors,which may be more readily available. The clock source on the HP router synchronous interface must be configured for "internal" when using these null-modem WAN cables.

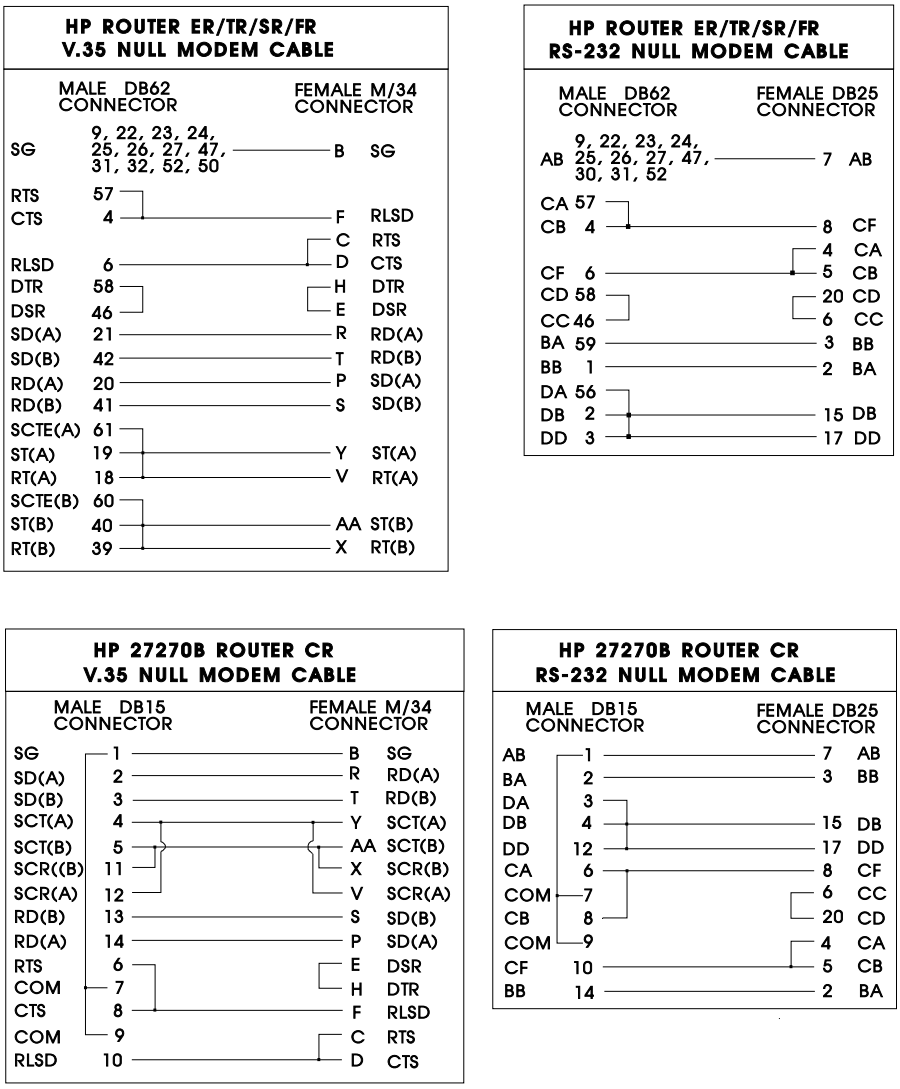

Figure 6. Pin-Outs For V.35 and RS-232/V.24/V.28 Null-Modem Cables for 62-pin Direct Connection

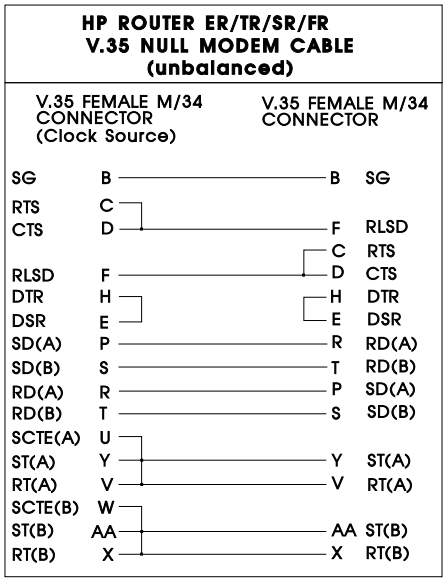

Figure 7. Pin-Outs for Null-Moden Cable for V.35 Connection

Special Considerations

Sync pass-through may be used only to provide communication between devices attached to synchronous ports on the router. Sync pass-through cannot be used to connect a synchronous device to a LAN-attached device.

In his June 21, 1991, *Data Communications* article “Router Backbones Unite Terminals and LANs”, Anura Gurugé describes a potential performance problem using sync pass-through. Extended LAN traffic is characteristically bursty. Bursts of LAN-to-LAN traffic can cause variable response-time delays on interactive terminals. According to Gurugé, variable response-time delays are one of the most annoying problems for interactive terminal users. To compensate for this delay, he recommends increasing backbone capacity by up to four times the sync pass-through line rate. Thus, up to 38.4 Kbit/s of additional capacity may needed on the backbone network for the addition of one 9.6 Kbit/s sync pass-through circuit.

Conclusion

Sync pass-through provides the opportunity to consolidate the traffic of many synchronous devices onto the corporate extended LAN internetwork. Consolidation of synchronous traffic can help reduce recurring monthly communication costs and costs associated with network support. Synchronous frames are transformed into LAN packets and bridged from source to destination. Synchronous devices must employ the HDLC data-link protocol or a derivative to transfer data using sync pass-through.

Routing with OSPF

Application Note

The capabilities of an internet are largely determined by its routing protocol. An internet's scalability, its ability to quickly route around failures, and the consumption of network resources by the routing machinery are all issues directly related to the routing protocol. With the release of HP router software revision 5.70, OSPF (Open Shortest Path First) is available in addition to RIP.

RIP (Routing Information Protocol) is probably the most widely used IP routing protocol. Its popularity stems from having been included with Berkeley UNIX (*Routed,* the routing daemon) and from being standardized by the IETF (Internet Engineering Task Force). RIP is documented in RFC (Request for Comments) 1058. RIP is a distance-vector protocol. A distance-vector protocol frequently (at 30-second intervals for RIP) sends its routing table (a vector of distances) to neighbor (adjacent) routers. When a router receives its neighbor's routing update, it compares the update with its own routing table and changes its routing table if necessary.

Distance-vector protocols are susceptible to two main problems. First, they can form routing loops, and second, they can be slow to converge. Convergence is the time required for the routing tables in all of the connected routers to stabilize after an event such as a network link failure. Routing loops and slow convergence are more likely to occur as network size and complexity increase. A fundamental assumption for routing (in the design of RIP) was that the internet contained at most a few hundred networks. Thus, RIP is not well suited for use with today's large corporate and government internets, which have thousands of networks. The limitations inherent in distance-vector protocols such as RIP and the lack of a standard routing protocol suitable for use in large internets are in large part responsible for the development of OSPF.

OSPF (Open Shortest Path First) is a new IP routing protocol. HP routers implement Version 2 of OSPF, which is documented in RFC 1247. Unlike RIP, OSPF employs a link-state algorithm (also referred to as a shortest-path-first (SPF) algorithm). Link-state algorithms are those in which each routing node floods information about its attached links to all other routing nodes.

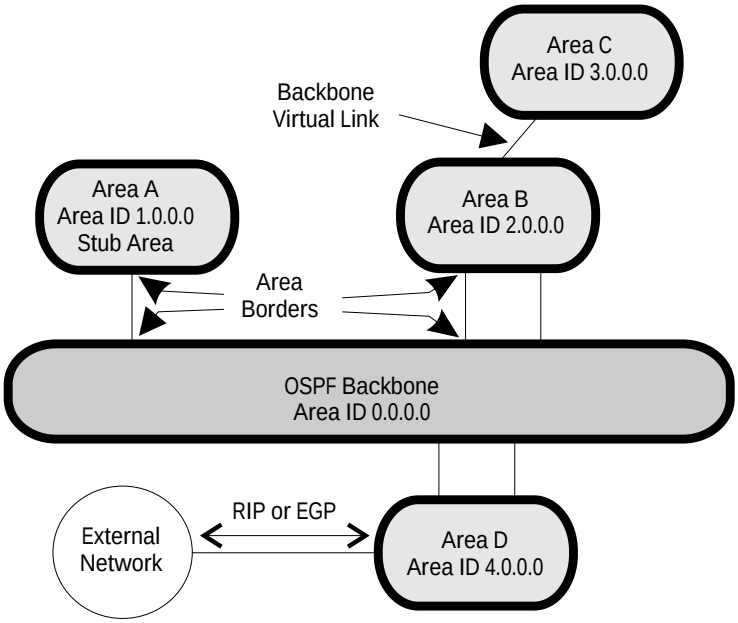

Figure 1. OSPF Autonomous System

Figure 1 shows an OSPF AS (autonomous system). AS is an IP term that refers to a collection of routers that all use the same IGP (interior gateway protocol). IGP is also an IP term. It refers to a routing protocol run within an AS. IGPs are a subset of routing protocols. They are referred to as IGPs to distinguish them from EGPs (exterior gateway protocols), another type of routing protocol. EGPs are used to route data between ASs. The Exterior Gateway Protocol (also known as EGP) is also the name of a specific EGP.

OSPF has features that:

- Improve routing effectiveness and efficiency.
- Conserve IP address space and increase addressing flexibility.
- Enhance network security.
- Increase routing flexibility.

These features make OSPF attractive for use on both small and medium-sized internets as well as large internets. This application note provides an overview of OSPF and describes selected features in more detail using model networks as examples.

Routing Improvements

OSPF has many features that are unarguably improvements over RIP. Those core features that improve the effectiveness and efficiency of routing include hierarchical routing, new routing metrics, the link-state protocol, and the topological database.

Hierarchical Routing—Areas

OSPF supports a routing hierarchy. Just as a hierarchical file structure is a better way to organize files than a flat file structure, so too a hierarchical network structure is a better way to organize networks than a flat network structure. The OSPF hierarchical structure helps to reduce the size of the topological database maintained by each router. (The topological database is discussed in more detail below.) It also helps to minimize routing control traffic. An AS may consist of one or more “areas”. Small networks may consist of a single area; large networks may contain many areas. Each area comprises a group of networks (or subnets). Each area in the AS is attached to the OSPF Backbone. An OSPF AS is logically a star: areas extend in a radial fashion from the Backbone. Note that a single-area AS would not have a Backbone area.

Data is routed within an area when the destination system is in the same area as the source. This means that when two systems inside area A (figure 1) communicate, the data is routed within area A. When the destination system is in a different area than that of the source, then data is routed from the source area to the Backbone area to the destination area. Therefore when a system in area A communicates with a system in area B, the traffic is routed from within area A to the Backbone. (The OSPF Backbone is an area in itself.) The router that connects area A to the Backbone is referred to as an “area border router”. The data is then routed through the Backbone to area B, where it is finally routed to the destination system.

Area and Router IDs. Both areas and routers have IDs (identifiers) that are used by OSPF to build its topological database. Both IDs are given in dotted decimal notation. This is the same notation used for IP addresses. The range of the identifiers is thus 0.0.0.0 to 255.255.255.255. Area identifier 0.0.0.0 is reserved for the Backbone area. Area and router IDs are not IP addresses, however, and thus such concepts as IP address classes, subnetting, broadcast addresses, etc., do not apply.

The OSPF standard does not provide guidance on the selection of area identifiers except the Backbone area (0.0.0.0). Several schemes have been proposed. A unified way to select both area IDs and router IDs is to assign IDs using a scheme compatible with that of the corporation or organization, such as `AREA.REGION.OFFICE.ROUTER`. In figure 1, area A has an area ID of 1.0.0.0, signifying a particular geographic area—a country, for example. Router IDs can then be selected based on their location within area 1.0.0.0. Thus the first router in office 1 of region 1 of area 1 receives the router ID 1.1.1.1. In any event, it is desirable to select identifiers that have some significance.

Area Sizing. A frequently asked question is: "How large should an OSPF area be?" Or: "What is the optimal number of routers in an OSPF area?" An area could contain as few as one router or as many as hundreds of routers. Carving a network up into smaller-sized areas will help to minimize the amount of OSPF protocol traffic. The costs associated with reducing the size of areas include additional complexity in the OSPF Backbone and the network as a whole. The answer to the question of area size probably has more to do with organizational structure than with the OSPF protocol. Generally, areas are suggested by the structure of the organization(s) responsible for managing the network. Within many large corporations, responsibility for network management is distributed. Corporate telecommunications departments may manage the network in and around the corporate offices or sites, while network management in other regions is handled by other groups in those regions. These organizational boundaries are also natural OSPF area boundaries.

Stub Areas. Stub areas are areas into which external routes are not propagated. The term external route has a particular significance in OSPF. Routing information provided to OSPF from any protocol other than OSPF itself is considered external. Thus routes provided by RIP or EGP as well as static routes are considered external. An OSPF router that interfaces to an external router is called an "AS Boundary Router." Consider an AS connected to the open internet through EGP. There are potentially thousands of routes for which reachability information could be obtained. To prevent this information from being propagated into an area, an area can be configured as a stub area. In figure 1, area A is a stub area. Therefore, external routing information received in area D will not be propagated into area A.

Virtual Links. In figure 1, area C is not directly attached to the Backbone. Instead, it is attached to area B through a "backbone virtual link". Virtual links allow areas to be configured where it would otherwise be inconvenient to do so due to the distance or cost to attach to the Backbone.

Metrics

The OSPF routing algorithm calculates the shortest path to each destination network. A path is composed of a series of network links from a router to a particular destination network. Each link is assigned a metric. The metric for a path is simply the sum of all the link metrics. The shortest path to a network is thus the path with the lowest metric. As networks become more complex, the result of alternate routes and varying link speeds, metrics become more important. OSPF provides a 16-bit (0 to 65535) dimensionless metric for the assignment of link costs and allows 24 bits for inter-area paths.

Table A. Metrics Based on Link Speed

Link Speed	Metric
100 Mbit/s	3
10 Mbit/s	10
4 Mbit/s	15
2.048 Mbit/s	32
1.544 Mbit/s	40
768 Kbit/s	75
512 Kbit/s	85
256 Kbit/s	95
128 Kbit/s	100
64 Kbit/s	105
56 Kbit/s	110
38.4 Kbit/s	120
19.2 Kbit/s	150
9.6 Kbit/s	200

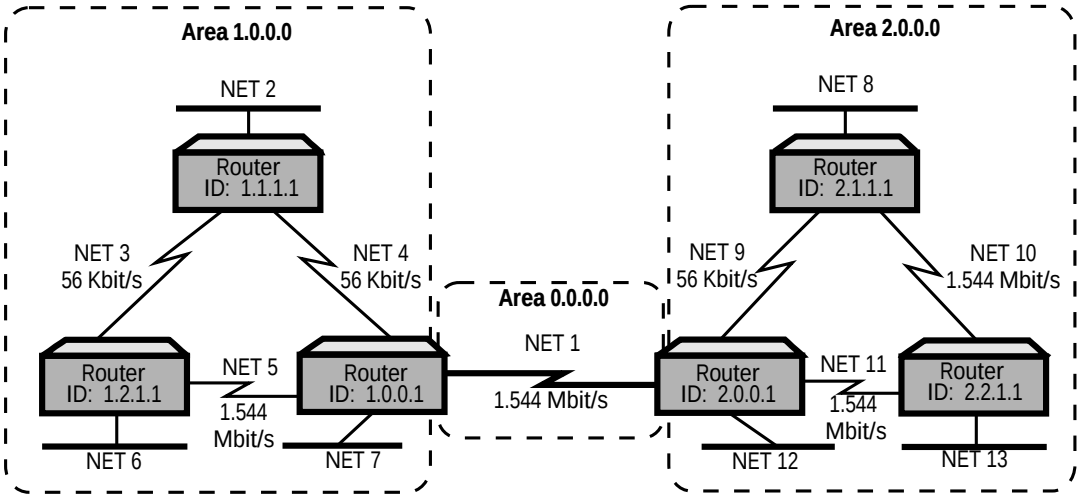

Figure 2. Multi-Area OSPF Internet

The OSPF standard does not provide guidance for metric selection or assignment. The easiest way to assign metrics is on the basis of link speed. Table A shows one possible scheme for selecting metrics based on link speed.

Each link in the network is assigned a cost (metric). In figure 2 the cost from router 1.1.1.1 to net 8 is the sum of the costs for net 4, net 1, net 11, net 10, and net 8. Assuming costs are assigned in accordance with table A, the cost of the route from router 1.1.1.1 to net 8 is thus 110 + 40 + 40 + 40 + 10 = 240.

Notice that the shortest path is not necessarily the one with the fewest number of hops, but rather the one with the lowest metric.

Link-State Protocol

When OSPF is started, usually during the router's boot procedure, it begins by synchronizing its database with those of its neighbor routers. Afterwards each router infrequently (at 30-minute intervals) floods LSAs (link-state advertisements) to all other routers in its area. Flooding is a way to send a message that will be relayed by all routers receiving the message. Received LSAs are used to build and maintain a topological database from which each router builds its routing table. There are several types of LSAs. Consider the network in figure 2. Each router in area 1.0.0.0 sends a router links LSA to all other routers in area 1.0.0.0 at 30-minute intervals or whenever a link state changes. The router links LSA includes the Router_ID of the router that originated the message (called the advertising router) and a description of each of the links connected to it. Link descriptions vary by the type of connected network. However, in this example the information in the LSA about a synchronous link includes:

- IP address of the link.
- Subnet mask used on the link.
- The router ID of the remote end router.
- The type of the link (point-to-point).
- Metric or cost assigned to the link.

Like the routers in area 1.0.0.0, routers in area 2.0.0.0 flood router-links advertisements to the other routers in area 2.0.0.0.

Routers maintain detailed topological information about area(s) of which they are members, and they maintain summary information about networks in other areas. Area border routers (1.0.0.1 and 2.0.0.1 in figure 2) exchange summary information about their own areas with other area border routers. Network summary information received by an area border router is then transmitted to routers in its attached area(s). The type of LSA used to send network summaries is called a summary-links advertisement. Like router-links advertisements, summary-links advertisements are also sent at 30-minute intervals and are triggered when a link state changes. A summary-links advertisement includes the following for each link advertised:

- IP network number or IP address of the link.
- Router ID (set to that of the local area border router).
- Subnet mask used on the link.
- Metric or cost from the area border router to the destination network.

Routers exchange LSAs to build and maintain their topological databases.

Topological Database

Routers in each area have identical routing databases. This topological database describes which routers are connected to which networks. Attributes of the networks and routers, such as subnet masks, metrics, etc., are also maintained in the database. Each router constructs its routing table from the database. The routing table contains the shortest path to every network the router can reach. The benefit of maintaining a topological database is that when a change occurs to the network, new loop-free routes can be computed quickly using minimal network resources. When a link outage occurs, for example, each router that detects the change floods LSAs that describe the change to all other routers in the network. Each router then modifies its database and recomputes the shortest path to each remaining network.

RIP, by comparison, maintains just the best route to any given network. When a change to the network occurs, routers must exchange their entire routing table with each neighboring router to relearn the best route to every network.

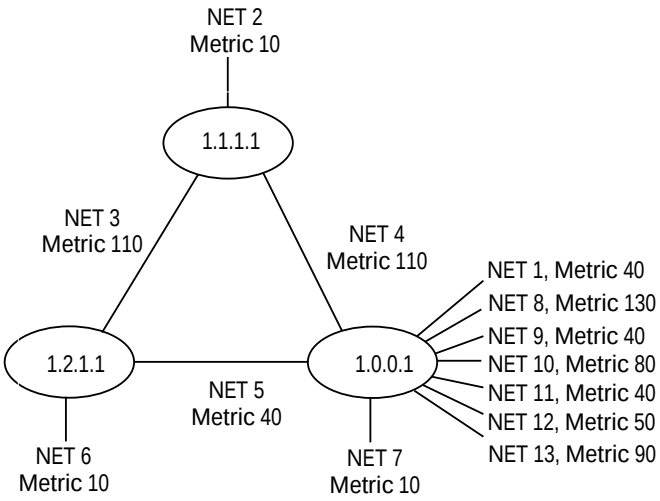

Figure 3. Area 1.0.0.0 Topological Database

Figure 3 shows a summary of the topological database of area 1.0.0.0 (from figure 2). Each router that is a member of area 1.0.0.0 has a copy of this database.

One aspect of the use of areas is that it simplifies the database. Notice that routers in area 0.0.0.0 and 2.0.0.0 do not appear in the topological database, although the networks in those areas do.

Conserving IP Address Space

On internets using RIP, the subnet mask used throughout the internet must be identical on all subnets. The information about individual networks included in a RIP update includes the network (or subnet) number and metric (hop count). RIP updates do not include the subnet mask associated with a network. This often results in the over-allocation of IP address space—especially on synchronous point-to-point links.

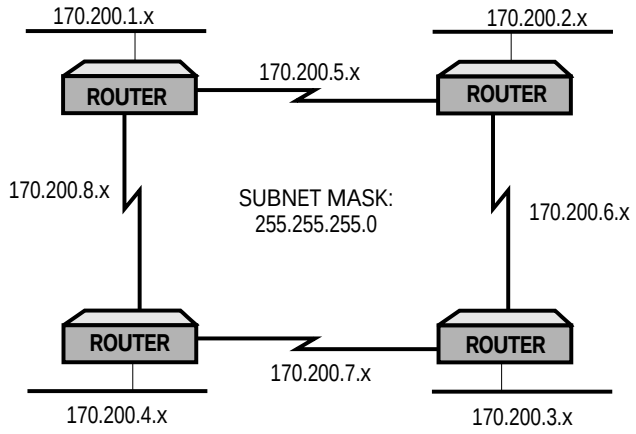

Figure 4. Subnetting in a RIP Internet

The internet in figure 4 uses the IP class B address 170.200.x.x. This internet is subnetted using the subnet mask 255.255.255.0. There are 254 possible addresses on each subnet (addresses 0 and 255 are reserved). On each of the wide-area network subnets, however, there are only two members, namely the router on each end of the point-to-point link. Of the 254 possible (or allocated) addresses, only two will be used. Thus 252 addresses are unused or wasted. Of the approximately 2000 IP addresses allocated in the internet in figure 4, roughly half are unusable because they are allocated to point-to-point networks. Variable-length subnet masks, a feature of OSPF, can be used to minimize the allocation of IP addresses, such as in the case of the point-to-point links in figure 4.

Variable-Length Subnet Masks

The mechanics of conserving IP address space using variable-length subnet masks are straightforward. Again consider the internet in figure 4. With OSPF the subnet mask can be specified per network or subnetwork. The subnet masks for the LAN subnets can be specified as before (255.255.255.0). This allows 254 addressable nodes per subnet. The point-to-point links, however, require only two IP addresses—one for each router. To restrict the number of addresses on the point-to-point links, a subnet mask is needed that allows only two addresses. It is tempting to use the subnet mask 255.255.255.254. The two addresses that this mask provides, however, are the “broadcast” address and the “any host” address for the subnet. Therefore, one additional bit is required in the subnet mask, so that reserved

addresses will not be used. Thus, the subnet mask 255.255.255.252 is the minimum subnet mask that will provide only two addresses per subnet. Now that a suitable subnet mask for use on point-to-point links has been determined, subnet numbers and addresses must be selected.

Table B. Point-to-Point Subnet and Address Definition

16 Bits	8 Bits	Upper 6 Bits	Lower 2 bits	32 Bits/32 Bits
Network Number	Subnet Number	Subnet Extension Component	Link Address Component	Full Point-to-Point Link Addresses
170.200	5	000000	01/10	170.200.5.1/170.200.5.2
170.200	5	000001	01/10	170.200.5.5/170.200.5.6
170.200	5	000010	01/10	170.200.5.9/170.200.5.10
170.200	5	000011	01/10	170.200.5.13/170.200.5.14

The following method of selecting subnet numbers and point-to-point link addresses illustrates one method of allocating addresses that seemingly maintains the address structure used in allocating LAN subnets. The first point-to-point subnet defined in figure 4 was subnet 170.200.5.x, the second was 170.200.6.x, and so on. To maintain consistency with that addressing scheme, the point-to-point subnets to be defined will use 170.200.5.x as the base for selecting link addresses. The last octet (the “x” octet) will be used to define individual point-to-point subnets and addresses. Table B shows the allocation of subnets and point-to-point addresses. Extending the subnet mask to include the upper 6 bits of the last octet for point-to-point links provides enough address space for 64 point-to-point links.

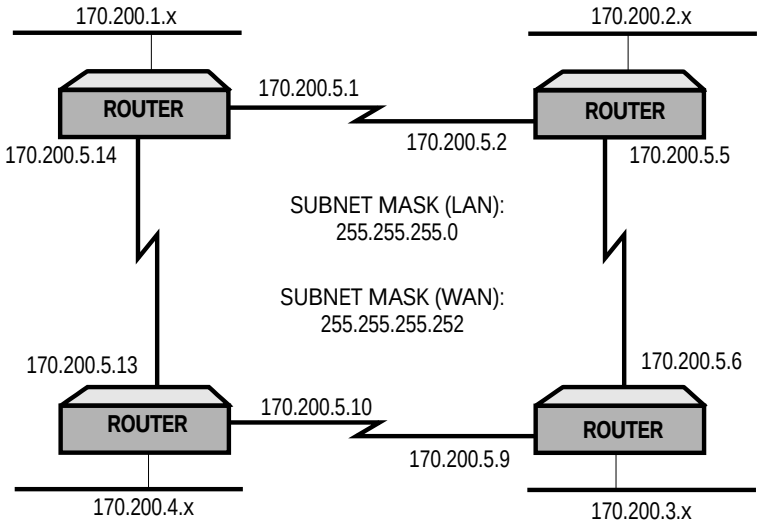

Figure 5. Internet with Variable-Length Subnet Masks

Using the addresses in table C, the internet from figure 4 is shown in figure 5. Full point-to-point link addresses are shown in figure 5 rather than subnets only, since the subnet mask is split on an octet boundary.

Reserved Addresses

When using variable-length subnet masks, special attention is required to avoid assigning addresses that are reserved. Avoiding the use of the “all hosts broadcast” address (all ones assigned to the host-address field of an IP address) and the “any host” address (all zeros assigned to the host-address field) was discussed above.

The other consideration that warrants special attention is avoiding the assignment of reserved subnet addresses. Reserved subnet addresses are those that are all ones (subnet broadcast) and those that are all zeros (any subnet). Variable-length subnet masks introduce extra complexity to this issue, since there may now be several different subnet definitions in a single network.

For example, consider the internet in figure 5. IP subnets in the range 170.200.0.4 through 170.200.0.248 appear to be valid IP subnets when used with the WAN subnet mask 255.255.255.252. However, these addresses are within the range of the reserved LAN subnet addresses, 170.200.0.0 through 170.200.0.255 (subnet mask 255.255.255.0), and thus must not be assigned. Similarly, subnets 170.200.255.4 through 170.200.255.248 appear to be valid IP subnets when used with the 255.255.255.252 subnet mask. This address

range is also within the reserved LAN subnet address range and must not be assigned.

The rule is to avoid assigning addresses within the reserved address ranges given by the subnet mask with the fewest number of bits in the subnet ID field. In figure 5 the subnet mask with the fewest number of bits in the subnet ID is 255.255.255.0. Thus the reserved address ranges are 170.200.0.0 through 170.200.0.255 and 170.200.255.0 through 170.200.255.255.

Addressing Flexibility

OSPF enhances IP addressing flexibility by allowing IP networks to be partitioned. Network partitioning is not permissible with RIP. A network becomes partitioned when one or more subnets of a network become separated from the other subnets of the same network by a second network. In figure 6, the class B network 128.1.x.x is partitioned by network 192.1.1.x.

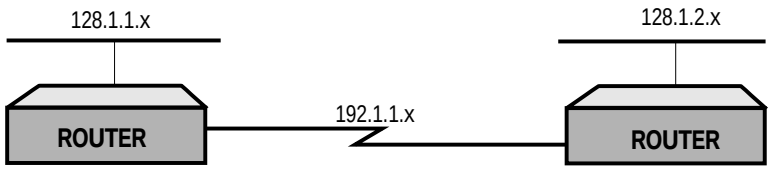

Figure 6. A Partitioned Network

To understand why network partitioning might be useful, refer again to figure 6. Suppose network 128.1.x.x is a large network with very few remaining subnets. Subnets of the class C network 192.1.1.x can be used for synchronous links so that none of the remaining class B subnets have to be used to extend the network.

Another situation in which support for partitioned networks is helpful is when networks are combined as a result of a merger or acquisition. Networks often overlap in this case, often resulting in capacity imbalances. Being able to partition a network allows the network's administrators much more flexibility in assembling the combined networks.

Network Security

Routing Authentication

To enhance security, routing updates may optionally be authenticated using a simple password. When routing authentication is enabled, all OSPF protocol packets are password protected. Passwords are from 1 to 8 characters and are configurable on a link basis. The determination to use routing authentication is made on an area basis. When routing authentication is used in an area, passwords must be configured on all links in the area.

Information Hiding

As discussed above, routers in each area have identical topological databases. Each router knows the topology of the areas to which it is attached. Only summary information is exchanged between areas. Thus, the topology of an area is hidden from routers outside the area.

Type-of-Service-Routing

Type-of-service routing is not yet available on HP routers. Currently, HP routers implement TOS 0 routing only. The following discussion is intended only to explain the type-of-service routing concept.

Type-of-service routing is based on the type of service defined in the Internet Protocol Specification, RFC 791. Three abstract quality-of-service parameters are given—delay, throughput, and reliability. These quality-of-service parameters (when used) are set in the IP protocol header by the system originating the IP datagram. Type-of-service values range from 0 to 7. Table C shows the four most common types of service. Note that the other types of service are based on combinations of the quality-of-service parameters given in table C.

Table C. IP Types of Service

Type of Service (TOS)	Delay	Throughput	Reliability	Service Description
0	0	0	0	Default
1	0	0	1	High Reliability
2	0	1	0	High Throughput
4	1	0	0	Low Delay

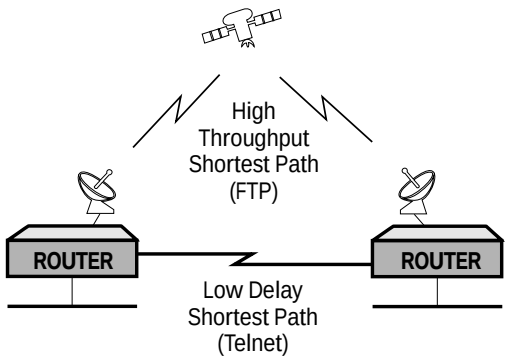

Figure 7. Routing Different Types of Service

When multiple types of service are supported by HP routers, routing decisions can be based on the TOS requested in the header of a datagram. Thus a separate set of routes can be calculated for each IP type of service. Conceptually, this will allow networks where file transfers (using FTP) can be routed over high-throughput routes such as satellite circuits, and time-sensitive data (using Telnet) can be routed over low-delay terrestrial circuits (see figure 7).

Equal-Cost Multipath

Equal-cost-multipath routing is not yet available on HP routers. The following discussion is intended only to explain the equal-cost-multipath routing concept.

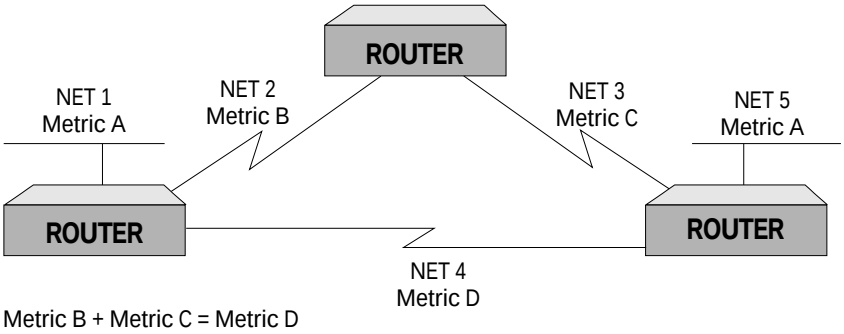

Figure 8. Equal-Cost Routes from Net 1 to Net 5

Consider the network in figure 8. There are two equal-cost paths from net 1 to net 5. The sum of the metrics for net 2 plus net 3 is identical to the metric for net 4. When packets destined for net 5 are received on net 1 by the router, packets will be forwarded on both paths. Obviously, careful attention must be paid to the assignment of link metrics to ensure that multiple paths are used to forward data when that is desired.

Conclusion

OSPF provides support for large networks using hierarchical network organization, improved metrics, and a link-state protocol. Additional features that help conserve IP address space, make addressing more flexible, and improve network security make OSPF attractive for use on many small and medium-sized internets as well.

OSPF features available in future router releases, such as type-of-service routing and equal-cost-multipath routing, will add even more value to an already valuable internetworking tool.

Application Note

Linking Up with Frame Relay

Overview

Frame Relay (FR) is a new wide area link technology. It was the first of several fast packet technologies to become commercially available. Other fast packet technologies include Switched Multimegabit Data Service (SMDS) and asynchronous transfer mode (ATM).

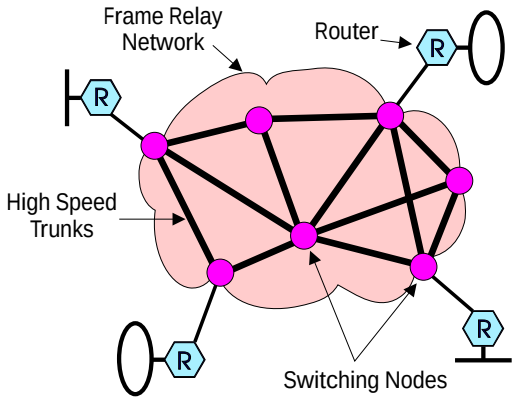

Figure 1. Frame Relay Network

A Frame Relay network is shown in figure 1. Frame Relay networks are composed of fast switching nodes connected by high-speed digital trunks. User devices, routers in this case, are connected to the switching nodes through digital circuits such as T1s. The Frame Relay standards define the interface between the user device and the Frame Relay network. How data is routed within the network is determined by the manufacturer of the

Frame Relay switches and/or the network provider. Frame Relay offers high performance, multiple access, and high reliability. This combination makes Frame Relay very well suited for LAN-to-LAN internetworking.

High performance or, rather, high throughput with low latency (delay introduced by the network), is achieved through a variety of factors. High-speed circuits are used for both Frame Relay access links as well as inter-switch trunks. Trunks in a Frame Relay network are typically very-high-speed digital links. In the U.S., T3 links (45 Mbit/s) are often used for this purpose.

With Frame Relay, data multiplexing occurs at the lowest possible layer—the data-link layer. Frame Relay standards do not specify a network layer. This reduces the complexity of Frame Relay implementations and helps to improve performance.

In contrast to X.25, Frame Relay networks do not provide error correction or flow control. Again, this reduces complexity and improves performance. Note that Frame Relay networks can detect errors. However, when frame errors occur, the frames in error are simply discarded by the network. The systems that originated the packets must retransmit the packets that were discarded (which typically occurs at the transport layer in a system's networking software).

Frame Relay allows for the transmission of large, LAN-sized frames (1600 bytes at minimum). This means that network software on both hosts and routers does not have to segment and reassemble packets. Frames may be up to 4500 bytes on some Frame Relay networks.

Frame Relay networks provide multiple access. Thus, many devices can attach to the same network and communicate through data-link layer addresses—similar to the way systems communicate on LANs. In contrast, private (point-to-point) leased lines offer access only at each end point.

High reliability is achieved through the use of reliable digital links, switching nodes with built-in redundancy, and a meshed network design where alternative data paths are available should a link or switch fail.

Frame Relay Physical Interface

Access to an Frame Relay network is usually provided through a digital circuit (local loop) from the Frame Relay carrier's point of presence (POP) to the customer premises. This circuit is terminated at the customer premises with a DSU/CSU. The physical interface on the DSU/CSU to which the router connects is a standard interface such as V.35, RS-449 (V.36), or X.21.

Frame Relay Connections

Figure 2 shows another view of the Frame Relay network. All of the routers are connected together in a fully meshed topology. Each router has a connection, a permanent virtual circuit (PVC), to each of the other routers in the network. (The PVC is illustrated by the lines in figure 2 that connect each of the routers inside the Frame Relay network.) Each connection is referenced by a Data Link Connection Identifier (DLCI). DLCIs are conventions between a router and a frame relay switch.

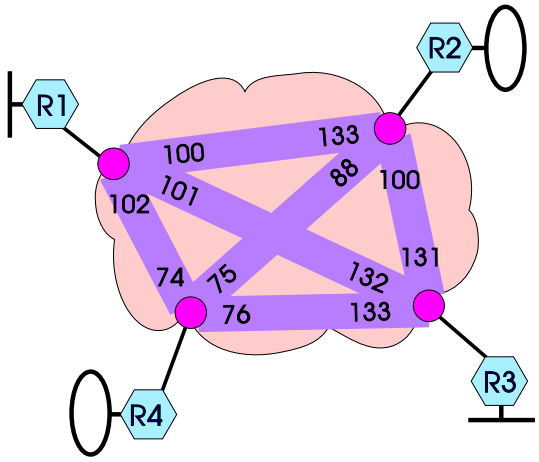

Figure 2. PVCs Connecting Routers

DLCIs are typically 10 bits wide. Were DLCIs unique throughout the network, the network would be limited to 1024 addresses—too few addresses for large public networks. However, DLCIs are not globally unique; instead, they are “locally significant”, which means they are unique only on the boundary Frame Relay switch. Thus, there are duplicate (or reused) DLCIs in a Frame Relay network. Also, routers on either end of a PVC typically have different DLCIs to reference the same connection. DLCIs 100 and 133 in figure 2 are examples of addresses that are reused.

The fully meshed topology shown in figure 2 is the topology that best exploits the value of Frame Relay. When compared to private-line pricing, a meshed topology using Frame Relay is often a fraction of the cost.

Ports and PVCs

Two important Frame Relay access characteristics are port speed and committed information rate (CIR). These characteristics govern the rate at which data may be transmitted into the network.

Port Speed Port speed is the maximum rate at which data will be accepted by the boundary Frame Relay switch to which a router is connected. This is the aggregate rate at which data can be transmitted to the network on all PVCs. Note that the port speed is frequently lower than the speed of the physical access link.

Committed Information Rate Committed information rate is the speed at which data is guaranteed to be accepted by the network on an individual PVC. The sum of the CIRs for all PVCs subscribed to is typically equal to the port speed. This is a general rule; it varies depending on the switches and software used to provide the Frame Relay networking service.

History

The motivations for Frame Relay began with X.25. The X.25 recommendation was released in 1974 by the CCITT. X.25 was widely accepted and implemented worldwide by the early 1980s. X.25 allowed a user to communicate with any other X.25 user worldwide simply by knowing the other party's X.121 address. (X.121 is a global data communication addressing standard.) In addition, X.25 pricing was usage based. X.25 was ideal for companies with many locations and with low to moderate communication requirements. X.25 could be implemented quickly, especially in situations that required international communications, and the costs were low when compared to private leased lines.

In the mid-1980s, many companies' communications requirements were increasing as a result of more powerful processors, more capable networking software, and new applications that required more bandwidth. It was becoming clear that X.25 would not be able to provide the throughput for future networking requirements.

Data communications standards committees were aware of these problems. New standards were being developed—the overall objective of which was to integrate both voice and data into one high-speed digital network. These standards—called ISDN (Integrated Services Digital Network)—are the basis upon which Frame Relay was defined.

The first formal Frame Relay specification was developed by the Frame Relay Forum (FRF), a group of companies that included Digital Equipment Corp., Cisco, Northern Telecom, and StrataCom. This specification, released in September, 1990, was based upon existing standards and ongoing work in the ANSI T1S1 committees. The FRF enhanced the basic data-transfer-protocol procedures defined in the ANSI standards, by adding the Local Management Interface (LMI) and some optional functions (which are discussed later).

In 1991, ANSI released T1.617, "ISDN—Signaling Specification for Frame Relay Bearer Service", and T1.618, "ISDN—Core Aspects of Frame Protocol for use with Frame Relay Bearer Service". In 1992, CCITT released Recommendation Q.922, "ISDN Data Link Layer Specification for Frame Mode Bearer Services", and Draft Recommendation Q.933, "Digital Subscriber Signaling System No. 1 (DSS 1), Signaling Specification for Frame Mode Basic Call Control".

Frame Relay services were first introduced in March, 1991, in the U.S. by WilTel, an interexchange carrier. Frame Relay announcements soon followed by other U.S. interexchange and local exchange carriers as well as the international exchange carriers.

Frame Relay Data-Link Interface

Table 1 shows the structure of frames transmitted by HP routers on Frame Relay networks. This frame structure was defined in the FRF's "Frame Relay Specification". The frame structure was based upon CCITT Recommendation Q.921, "LAPD", and from work being done in the ISDN area and in the ANSI committee T1S1. It was then standardized by ANSI in Standard T1.618 and by the CCITT in Recommendation Q.922. Q.922 is also referred to as "Link Access Procedure to Frame Mode Bearer Service (LAPF)". In this document, we will refer to the Frame Relay frame structure as Q.922.

Table 1. Q.922 Frame Structure

	8	7	6	5	4	3	2	1
Octet 1	**Flag (7E hex)**							
Octet 2	**Q.922 Address**							
Octet 3	**Q.922 Address**							
Octet 4	**Control (UI=03 hex)**							
Octet 5	**Optional Pads (00 hex)**							
Octet 6	**NLPID**							
Octet 7 Octet N-3	**Data**							
Octet N-2	**Frame Check Sequence**							
Octet N-1	**Frame Check Sequence**							
Octet N	**Flag (7E hex)**							

A standardized procedure for transmitting multiprotocol data over Frame Relay was defined in Internet RFC 1294, "Multiprotocol Interconnect over Frame Relay". The shaded fields in table 1 are defined by RFC 1294. RFC 1294 defines standard procedures for sending multiprotocol (connectionless network-layer) data over Frame Relay networks. All popular LAN protocols such as TCP/IP, Novell IPX, and AppleTalk have connectionless network layers. RFC 1294 does not address interconnecting connection-oriented protocols such as IBM's SNA over Frame Relay.

Like LAPB (the link-access procedure used on X.25 networks) and LAPD (the link-access procedure used in ISDN networks), Q.922 frames are delimited by a flag character (7E hex), and for data integrity, a frame check sequence (CCITT-CRC) is appended to the user data.

Table 2. CCITT Rec. Q.922 Two-Octet Address Field Format

8	7	6	5	4	3	2	1
Upper DLCI						C/R	EA
Lower DLCI				FECN	BECN	DE	EA

Q.922 Address

The format of the Q.922 address field is shown in table 2. In addition to the two-octet (byte) format shown in table 2, there are three- and four-octet formats. These formats can be used to provide additional addressing capability and will be used in the future to provide switched virtual circuits. For now we will consider only the two-octet format, since it is typically the only addressing format supported on today's public Frame Relay networks.

Data-Link Connection Identifier The six-bit upper DLCI and four-bit lower DLCI are concatenated to form a ten-bit DLCI. This provides for up to 1024 DLCIs per Frame Relay switching node (again, considering only the two-octet address format).

Control Field and Command/Response Bit Frame Relay frames are sent with the control field (octet 4 in table 1) set to 03 hex. This encoding indicates that the frame is a UI (unnumbered information) type of frame (identical to the encoding of frames transmitted on most LANs). The command/response bit (denoted by "C/R" in table 2) is always encoded with a 0, which indicates that the frame is a command. (UI frames are command frames.)

Address Extension Bits The EA bit in the Q.922 address field is the address extension bit. The EA bit in the last octet of the Q.922 address field is always set to 1, which indicates that it is the last octet of the address field. All preceding octets in the address field have the EA bit cleared.

Congestion Control The most interesting aspects of the Q.922 address field are the three bits used for congestion control—FECN, BECN, and DE.

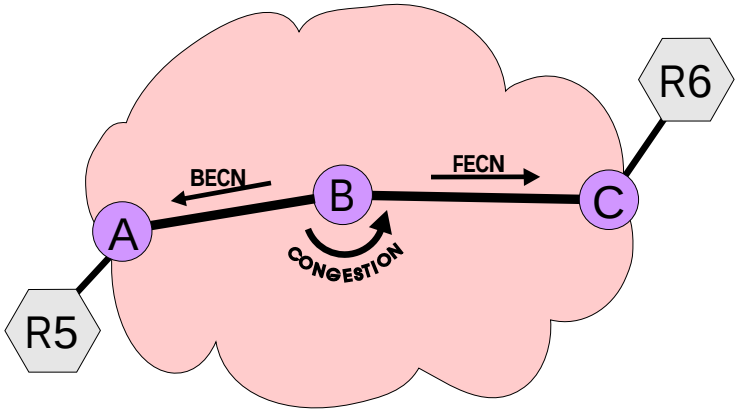

Figure 3. Congestion on a Frame Relay PVC

Figure 3 shows a Frame Relay network with two attached routers. Router 5 (R5) is sending data to R6 on a particular PVC. Assume that node B in the Frame Relay network detects that the network is congested or about to become congested on the route from R5 to R6. Node B responds to this network congestion condition by setting the Forward explicit congestion notification (FECN) bit in frames going from R5 to R6. The setting of this bit tells R6 that incoming data on the PVC from R5 is experiencing congestion.

Alternatively, node B may discard the frame from R5 if the discard eligibility (DE) bit in the frame is set. The discard eligibility bit would typically have been set by node A in frames that exceeded the CIR of the PVC from R5 to R6. HP routers never set the DE bit in frames transmitted to the Frame Relay network.

To notify R5 that a congestion condition exists on the route to R6, node B sets the backward explicit congestion notification (BECN) bit in frames going from R6 to R5 (same PVC as before). This signal is intended to signal R5 to reduce its transmission rate on that PVC.

Frame Relay standards do not require attached devices (routers) to respond to congestion notifications. In fact, HP routers ignore these indications, as do most other router implementations, although the number of BECNs and FECNs received is maintained in the MIB.

Congestion notifications are ignored for a couple of reasons. First, routers do not have good mechanisms for notifying end systems on attached LANs to slow their transmissions to a particular destination. The TCP/IP and OSI transport layers, however, have implicit congestion avoidance mechanisms. When these protocols begin to detect network delays, they reduce their transmission windows on the affected sockets or SAPs and thereby help relieve congestion conditions.

Second, since routers are not required to respond to congestion notifications, there is an issue of fairness. If more than 1 router (different vendors) is attached to a node in a Frame Relay network and one behaves responsibly while the other does not, the one behaving responsibly is penalized by realizing lower throughput rates on the affected PVCs.

Optional Pads

Pad characters may optionally be included in the Frame Relay header. Pads allow the frame to be extended to an even or odd number of bytes or words in length to suit the needs of the frame's transmitter.

NLPID

The Network Layer Protocol Identifier (NLPID) identifies the type of network on which the packet was generated, for example, TCP/IP, Novell IPX, AppleTalk, etc. NLPID values are defined in ISO 9577. Of the routable protocols supported by HP routers, only Internet Protocol (the IP in TCP/IP) has a defined NLPID field. Therefore, a couple of additional fields may be required to encode the network type. Additional fields may also be added if the packet is to be bridged. Table 3 shows the encoding of the NLPID field.

Table 3. NLPID Field Encoding

00 (hex)	**Not Used**
80 (hex)	**SNAP**
81 (hex)	**ISO CLNP**
82 (hex)	**ISO ES to IS**
83 (hex)	**ISO IS to IS**
CC (hex)	**Internet Protocol (IP)**
CE (hex)	**EtherType**

When a packet other than TCP/IP, such as Novell IPX, must be transmitted, RFC 1294 specifies that the NLPID field is encoded with 80 hex, which indicates that a Subnetwork Access Protocol (SNAP) header follows. The SNAP header and additional data will be appended to the NLPID field and used to encode which of the other routable protocols the Frame Relay frame contains. Table 4 shows the encoding of a Novell IPX frame.

The encoding of the three-octet SNAP header is 00 00 00 hex. This indicates that an Ethernet type field follows. 8137 hex is the Ethernet type field for Novell IPX. The Novell header and data follow in the remainder of the Frame Relay frame.

Table 4. Routing Novell IPX — NLPID Continuation

00 (hex)	**SNAP Header Indicating EtherType Follows**
00 (hex)	
00 (hex)	
81 (hex)	**EtherType (Novell)**
37 (hex)	
● ● ● ● ●	**Novell Header plus Data**

Bridging Over Frame Relay

The encoding of the NLPID field for bridged packets is the same as that used for network layers not having a specific NLPID value—80 hex. This again specifies a SNAP header; however, the encoding of the SNAP header is 00 80 C2, the organizational identifier of the IEEE 802.1 Committee. Table 5 shows how a bridged frame is encoded.

Table 5. SNAP Header for Bridging

00 (hex) 80 (hex) C2 (hex)	**IEEE 802.1 Organization Code**
xx (hex) xx (hex)	**Protocol ID PID**
● ● ● ● ●	**Complete Bridged Frame**

The PID (Protocol ID) field identifies the media type and whether or not the LAN FCS is appended to the bridged frame. The PID values are shown in table 6.

Table 6. PID Values for Bridged Frames

With Preserved FCS	Without Preserved FCS	Medium
00 - 01	00 - 07	802.3/Ethernet
00 - 02	00 - 08	802.4 (Token Bus)
00 - 03	00 - 09	802.5 Token Ring
00 - 04	00 - 0A	FDDI
00 - 05	00 - 0B	802.6
00 - 0E		**Bridge BPDU**

At this point we have looked at Q.922, the Frame Relay data-link interface. We have seen how multiprotocol user data is encoded and exchanged in a Frame Relay network. In the next section we will see how the Frame Relay link and each PVC is activated (or deactivated) and how the status of each link and the associated PVCs is maintained. These are all aspects of the LMI (Local Management Interface). We will look at an optional extension to the LMI for handling LAN-originated multicast/broadcast frames. Additionally, we will briefly discuss Internet RFC 1293 (Inverse ARP) and issues related to address resolution.

Local Management Interface

Local Management Interface is the phrase used in the FRF's "Frame Relay Specification" to describe the mechanisms for providing a user device with network configuration and status information. Both ANSI T1.617 Annex D and CCITT Q.933 Annex A have the same capability for providing configuration and status information for PVCs as the LMI has. These three standards are collectively referred to here as LMI unless otherwise indicated. It should be noted that HP routers do not yet support CCITT Q.933 Annex A.

LMI procedures include:

- Notification of the addition or deletion of a PVC.
- Notification of the availability of a PVC.
- Link integrity verification (keep alive).

PVC Status Messages

There are two defined LMI message types: Status and Status Enquiry. The Status message is always sent to the attached router by the network in response to a Status Enquiry message. A Status Enquiry message is sent by a router at a predetermined polling interval to request status from the network.

There are two types of Status Enquiry/Status messages exchanged by routers and the Frame Relay network:

- Link Integrity Verification (Keep Alive). This message is sent periodically to maintain link integrity, or in other words, keep the link alive. Routers and the network exchange their send and receive sequence numbers to maintain synchronization.
- Full Status. This message provides routers with the status of all available PVCs. Additionally, send and receive sequence numbers are exchanged as in the case of the keep alive message.

Periodic Polling Routers are required to periodically poll the network with a Status Enquiry message. This requirement holds whether data is being exchanged over the interface or not. The default polling interval is ten seconds. A number of keep alive polls are sent in between full status polls.

Link Integrity Verification (Keep Alive) Both the router and the network verify the integrity of the link with every Status Enquiry/Status exchange. Send and receive sequence numbers are included in each LMI message exchanged with the network. The network and attached routers verify that LMI messages are sent and received at the proper intervals and ensure that sequence numbers increment properly. If an LMI message is not received within a predetermined time period or the sequence numbers are not the expected values, an error occurs. Routers and the network allow a configurable number of errors within a specified time interval (window). If this value is exceeded, the link is removed from service.

This sequence of events is shown in figure 4. Both the “poll interval” and the “interval between full polls” are configurable parameters, the values of which should be set in accordance with the network provider’s subscription parameters.

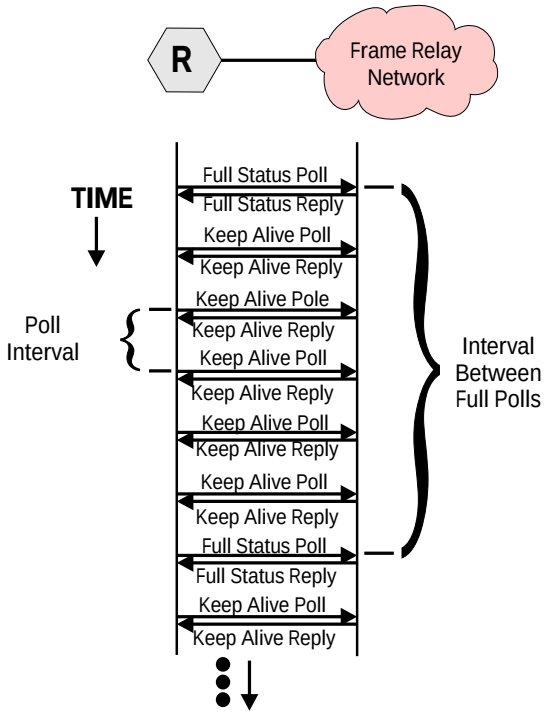

Figure 4. Interleaving of Full Status and Keep Alive Polls

Full Status When a router attached to a Frame Relay network is initialized, the first message that the router sends to the network is a full status poll. The network responds with a full status reply. The status information returned by the network indicates which PVCs are available for use. The full status reply lists each of the DLCIs available and provides two bits of status information per DLCI. The first status bit is called the “New” bit; the second status bit is called the “Active” bit. Table 7 shows how these bits are used.

Table 7. DLCI Status bits returned in "Status" message

Bit	Value	Description
New	0	PVC is already present
New	1	PVC is new
Active	0	PVC is inactive (PVC is deleted)
Acitve	1	PVC is active

Figure 5 shows four routers attached to a Frame Relay network. Each router has a PVC to each of the other routers on the network. Assume that R1 is the first to initialize (power on) and become active on the network. The first message that R1 sends to the Frame Relay network is a full status poll. This indicates to the network that the router and the link to the router are now functioning. The network returns a full status reply. This status reply, however, contains an empty list of DLCIs (no DLCIs and accompanying status information). This is because none of the other routers are yet active.

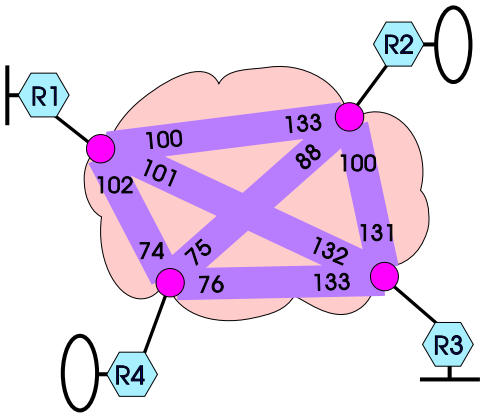

Figure 5. PVCs Connecting Routers

Now, let's suppose that R2 powers on, and its Frame Relay interface is initialized and enabled. Like R1, the first message it sends to the network is a full status poll. As before, the network responds with a full status reply. This time, however, the list of DLCIs is not empty. The full status reply contains DLCI 133 and indicates that it is both new and active (see table 7). This indicates to R2 that it may now communicate using DLCI 133 (to R1). At this instant, however, R1 still has no active DLCIs and is not able to communicate with R2. This situation is resolved after R1 sends its next full status poll

to the network. The network will respond with a full status reply, which indicates that DLCI 100 is available (active and new bits set). Both routers are now able to communicate.

When R3 comes on line and sends its full status poll to the network, the network will respond with a full status reply indicating that DLCIs 131 and 132 are available. R1 and R2 will learn about the availability of PVCs to R3 after they send their next full status poll to the network. This continues until each router knows about the PVCs to each of the other routers.

If one of the routers fails or is powered off, the remaining routers are informed that its DLCI to the affected router is “inactive” (active bit cleared).

Multicast

An optional feature of the FRF’s LMI is the ability to send a frame to all user devices that belong to a multicast group. In essence, a router transmits a single frame to the network (on a special multicast DLCI), and that frame is delivered to all members of the multicast group. This is a useful feature for the transmission of routing updates and other types of broadcast messages. The FRF’s LMI reserved four DLCIs for this use. The multicast feature is not available in ANSI T1.617 Annex D.

Without multicast support, by contrast, whenever a broadcast message such as a routing update must be sent, it is transmitted on each DLCI. In figure 2, for example, when R1 broadcasts a routing update (RIP, OSPF, etc.), it is sent three times—once on each DLCI to each of the other routers.

As a practical matter, the multicast feature is of limited value. Since there are only four DLCIs reserved for multicast, it is impractical to offer this feature on public networks. Multicast support probably offers the greatest utility on private Frame Relay networks.

StrataCom, one of the leading suppliers of Frame Relay switches to public networks, has recently withdrawn its support for multicast capability on its Frame Relay switches. It is therefore doubtful that this feature in its current form will find much use on today’s Frame Relay networks.

Address Resolution

Internet RFC 1293, “Inverse Address Resolution Protocol”, describes additions to ARP (Address Resolution Protocol) that are intended to reduce the amount of address resolution traffic in a Frame Relay network. An ARP packet is sent whenever a TCP/IP system wants to communicate with a system it hasn't communicated with recently. The sender of the ARP needs to determine the hardware address of the destination system. Normally this packet is broadcast to all machines on a network.

Inverse ARP reduces address resolution traffic by directing an Inverse ARP packet to each system when its DLCI becomes active. What Inverse ARP does is to match up an unknown IP address with a known hardware address (DLCI). When enabling Inverse ARP on HP routers, it is a good idea to disable the ARP Cache Timer of the Frame Relay circuit in the network interface definition using the router's Configuration Editor.

Protocol Analysis/Testing

The HP PT502 and HP PT302 are high-performance wide-area network protocol testers, which provide powerful monitoring, analysis, emulation, and conformance testing capabilities for a variety of protocols including Frame Relay at speeds up to 2 Mbit/s. The PT502 is a multiport protocol tester designed for lab use and high-level network maintenance. The PT302 is a companion of the PT502; it has a single test port and is well suited for field use.

These protocol testers may be used to:

- Monitor Frame Relay protocol between a user device and the Frame Relay network. Triggers, filters, and detailed decoded displays reduce troubleshooting time. All identified protocol violations are highlighted in red, allowing the user to identify the problem quickly.
- Perform statistical analysis of network performance. Statistical analysis is useful for predicting network usage versus capacity, predicting growth trends, and preventing network overload conditions.

- Emulate either user device or Frame Relay network. User device/ network emulation can be used to test or troubleshoot particular capabilities of an interface. The “Test Manager” in the emulation package is a comprehensive environment for automating the entire array of Frame Relay test functions.
- Generate traffic (load) to the device under test. Load generation can be used to stress test Frame Relay switching equipment and observe the behavior of the device under extreme load condition.
- Test the protocol implementation for conformance to ISO 9646 recommendations. The test report of each test case outlines the test result along with detailed diagnostic information, protocol traces, and problem identification. This information can be used to help eliminate design or implementation flaws before the development of the equipment is completed, thus reducing the time to market as well as costs.

The product numbers for the PT502 and PT302 protocol testers and associated Frame Relay software packages are:
- PT502: HP E3910B
- PT302: HP E3939B
- Frame Relay Analysis software package: HP E3946A
- Frame Relay Emulation software package: HP E3947A
- Frame Relay Common Control software package: HP E4089A
- Frame Relay ACT-FR Test Suite software package: HP E4092A

HP Implementation—Summary

The Frame Relay interface implemented on HP routers is summarized below.

Protocols Supported over Frame Relay:
- IP, Novell IPX, AppleTalk Phase 2, DECnet IV, XNS, and bridging.

Protocol Encapsulation Method:
- Internet RFC 1294

Address Resolution:
- Support for Inverse ARP (RFC 1293) may be enabled or disabled.

Data-Link Layer Specification:
- CCITT Rec. Q.922 (two-, three-, and four-byte addressing formats). Compatible with Frame Relay Forum's "Frame Relay Specification" and ANSI T1.618.
- CCITT Rec. Q.922, "PreDraft Standard" and the earlier March and November versions.
- CCITT Rec. Q.921.

Maximum Frame Size:
- 1600 Bytes (Ethernet/802.3)
- 4500 Bytes (Token Ring/802.5)

Management Interfaces:
- Frame Relay Forum LMI
- ANSI T1.617 Annex D
- Null Management Interface. (PVCs must be statically configured.)
- LMI Switch and ANSI T1.617 Annex D Switch. These interfaces are included for test purposes.
- CCITT Rec 933 Annex A (not yet supported)

Multicast Support:
- Frame Relay Forum LMI Multicast Extension

Flow Control:
- LMI XON/XOFF

DE Bit Setting:
- The DE bit is never set.

Response to FECN/BECN:
- Received FECNs and BECNs are counted. No action is taken to avoid congestion.

Monitoring of Committed Information Rate (CIR):
- None.

Frame Relay MIB:

The Frame Relay information base contains the following three tables and associated entries:

Data-Link-Connection Management Interface Table (1 DLCMI table per physical circuit)

1. LMI in use (FRF LMI, ANSI Annex D)
2. Address Format (Q.921, Q.922, etc.)
3. Address Length (2, 3, 4 octets)
4. Maximum number of PVCs supported
5. Poll Interval
6. Interval Between Full Polls
7. Monitored Events
8. Error Threshold

Circuit Table (1 per physical circuit)

1. DLCI #1 (first table entry)
   - DLCI # (actual DLCI used)
   - DLCI creation time
   - Circuit state
   - Time since last circuit state change
   - Frames received
   - Frames sent
   - Octets received
   - Octets (bytes) sent
   - BECNs received
   - FECNs received
2. DLCI #n (next table entry, etc.)

Error Table (1 per physical circuit)

1. Error data (string containing packet in error)
2. Error time (time the error was detected)
3. Error type

Case Study

Frame Relay Network Design: Fleet Call, Inc.

Wendy Pinos, Network Consultant
Hewlett-Packard Company

Company Overview

Fleet Call, Inc., is a rapidly-expanding wireless communication company, located in the major market areas across the United States. Fleet Call is currently the nation's second largest provider of specialized mobile radio services. Starting August 1993, Fleet Call intends to provides high-quality all-digital mobile communication services such as mobile telephone, dispatch, paging, and data services to customers. In February of 1991, Fleet Call received authorization from the FCC to build Digital Mobile networks. This will allow Fleet Call to begin the process of building a single, nationwide, all-digital network, using compatible technologies and providing compatible services across all geographic areas. The move to digital technology will provide enhanced service and superior operation for mobile radio customers.

Fleet Call plans to commence its operation of the advanced Digital Mobile network, beginning in Los Angeles in August 1993,
San Francisco, New York, and Chicago in 1994, and Dallas/Fort Worth and Houston in 1995.

Business Need

To support the business services, Fleet Call is installing a series of HP 9000 UNIX systems. These systems will run subscriber maintenance and financial applications critical for Fleet Call's day-to-day operations. The HP 9000 systems are housed in a centralized data center.

All remote Fleet Call sites will need real-time access to the host systems. To achieve this connectivity, Fleet Call has designed a robust and high-performing wide area and local area topology.

Applications

Fleet Call's host systems will run several applications. In addition to the subscriber and financial applications, a number of engineering applications are also under investigation. Users will access the host applications from PC and Macintosh workstations distributed across the network. The workstations will be equipped with network software and hardware to allow Telnet access to the host systems over TCP/IP for terminal emulation. Host printing will also be spooled over the network to TCP/IP-equipped HP LaserJet printers. It is Fleet Call's goal to provide a completely transparent network, with 100% uptime.

Network Topology

Fleet Call's network focuses principally at providing connectivity to the HP 9000 systems. A two-tiered network topology is used: First, a network backbone interconnects a key site in each of the principal market areas. This backbone will provide a "data highway", allowing data communication between the remote sites and the data center at high bandwidths.

At the second tier, each of the backbone sites concentrate wide area connections from the remote offices located in that market area. These nearby offices or "tail sites" will rely on the backbone connection for communication to other market areas or the data center. This topology is summarized in figure 1.

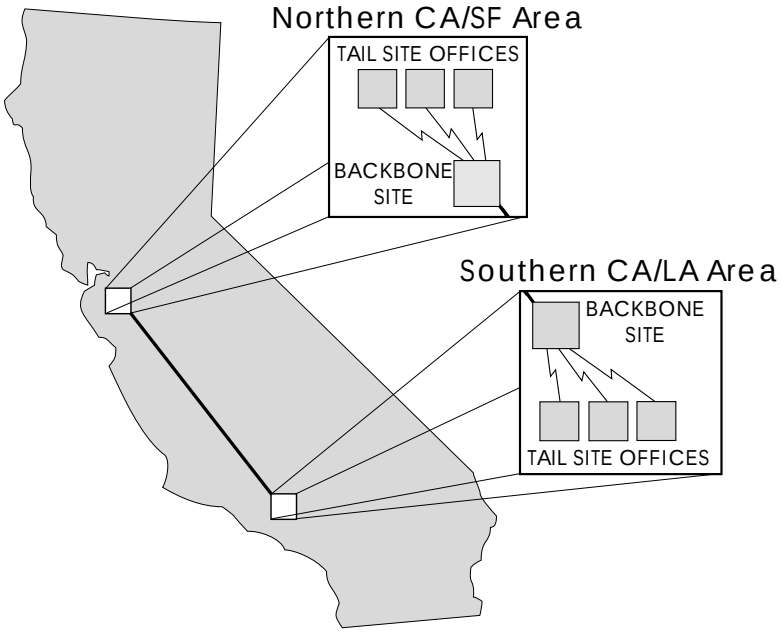

Figure 1. Two-Tiered Topology

Each tail site will be equipped with a small network router that allows connection of the local PC network to the backbone through wide area links. Routers were chosen rather than bridges because routing technology provides a greater level of network traffic filtering. The Fleet Call network will grow to incorporate many hundreds of nodes, and optimizing the bandwidth utilization of the wide area links is essential to preserve high performance and minimize cost.

The backbone sites are equipped with a larger network router that concentrates the lines from all the tail sites and also provides a connection to the backbone network. The backbone router will also concentrate the local LAN subnets within the building itself.

This network design offers several advantages. By using routers at each remote site, the network traffic will be filtered and localized to the specific workgroups. Two levels of filtering are used: first, at the local site itself, and second, at the backbone. Thus, any network traffic with destinations within a particular office, or within offices in that specific market area, will be filtered from transmission onto the backbone. This preserves the backbone bandwidth for essential traffic.

Second, using the backbone/tail-site approach, the cost of the wide area links can be optimized. The tail-site links will fall principally within a LATA (Local Access and Transport Area—the service area) boundary, allowing a local operating company, such as Pacific Bell, to provide the link service. The link costs are usually based on distance between connection points, and the short intra-LATA links from the backbone to the tail sites will be much more cost effective than providing multiple inter-LATA links directly to the data center on the backbone. Since network traffic will be concentrated at the backbone router, Fleet Call can take advantage of lower-bandwidth lines out to the tail sites and higher-bandwidth lines on the backbone itself.

In evaluating an appropriate backbone technology for Fleet Call, the following basic requirements were considered:

- high-bandwidth transmission media
- built-in link redundancy
- easy growth and reconfiguration
- cost-effective operation

The traditional design for wide area networks uses point-to-point dedicated circuits. Dedicated circuits provide guaranteed bandwidth with high performance. With point-to-point circuits, Fleet Call could still implement the two-tiered network described above. However, the point-to-point circuits offer only a single circuit on the backbone, so Fleet Call would need to implement a meshed network of multiple point-to-point circuits to achieve any kind of link redundancy. A sample mesh network is shown in figure 2.

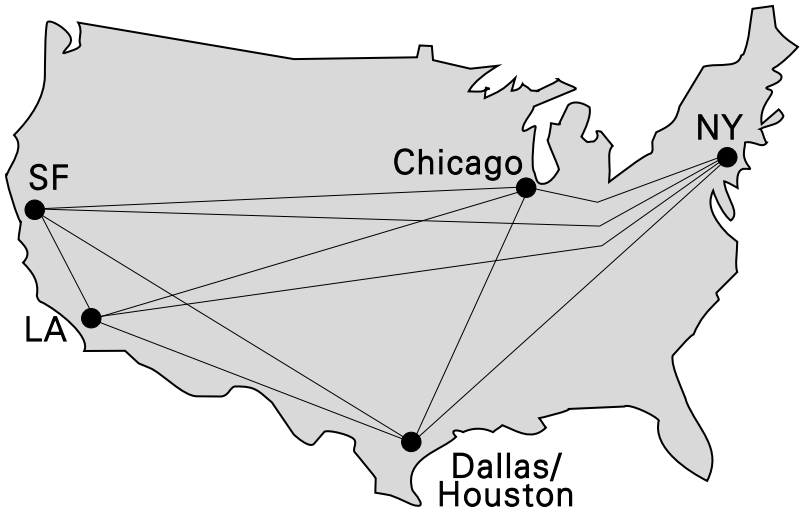

Figure 2. Point-to-Point Meshed Network

Circuit costs are based on distance, and Fleet Call's wide geographic dispersion would mean sizable monthly circuit costs to effectively run a meshed network between the backbone sites. In addition, Fleet Call would be solely responsible for the troubleshooting and problem isolation in a point-to-point circuit network. This requires a higher level of technical sophistication and knowledge from the Fleet Call network staff.

A number of alternative backbone technologies are emerging in the industry. Long-distance carriers have made the most progress in providing frame relay services. Frame relay bases its design on packet-switching technology, as does X.25. However, unlike X.25, frame relay provides its control through only OSI layers 1 and 2, which preserves network performance. The design of a frame relay network provides an fully meshed network that compensates for point failures with automatic rerouting. Customers have a single point of entry to the frame relay network, and all network routing happens transparently. This is depicted in figure 3.

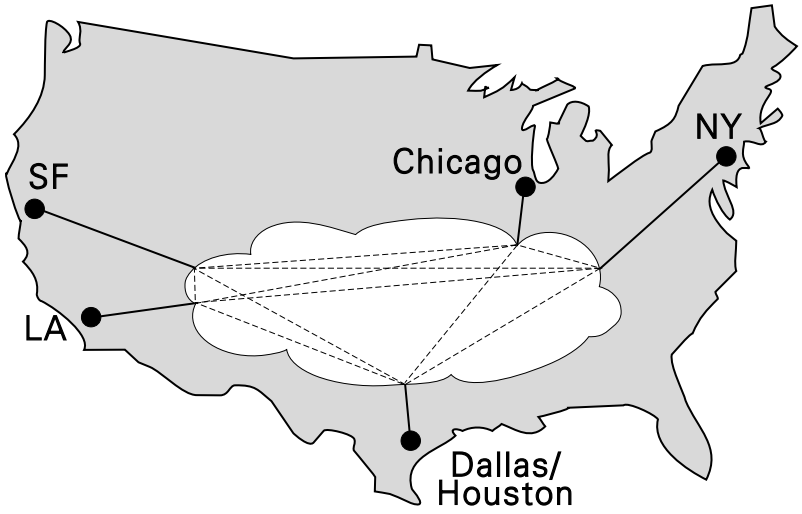

Figure 3. Frame Relay Network

Frame relay is also very suitable for bursty data networks, where irregular traffic patterns cause variable bandwidth demands. A frame relay backbone network would provide Fleet Call with built-in redundancy and high performance. Compared to a meshed topology with dedicated circuits, frame relay is more cost-effective. In addition, the selected long-distance carrier assumes total responsibility for failures within the meshed network, thereby relieving Fleet Call of this responsibility.

Both the dedicated circuit and frame relay approaches have merit. For the final wide area topology, Fleet Call selected a combination of these solutions. For the high-speed network backbone that provides mission-critical connections, a frame relay network is used. For links to the tail sites, point-to-point dedicated circuits running back to the backbone are suitable. The wide area design is shown in figure 4.

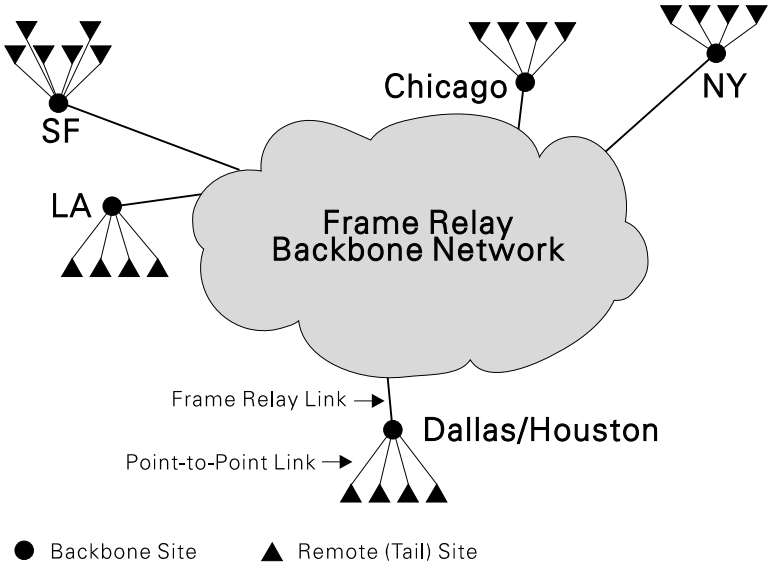

Figure 4. Fleet Call WAN Topology

Each office site is equipped with 10Base-T networks running Novell NetWare to provide local file and disk sharing. Except for electronic mail transfer, the NetWare traffic is localized within each office, preserving the bandwidth of the WAN for the TCP/IP host-directed traffic.

Fleet Call will use the HP OpenView Network Management system. It allows Fleet Call to monitor and manage all the network devices that have IP addresses or can be queried with SNMP. This includes the routers, network hubs, print spoolers, PCs running the PC/TCP software, UNIX workstations, and Macintoshes running TCP software.

Performance

As of mid-1993, Fleet Call continues to roll out the network to remote sites. Performance will depend on the bandwidth of the remote links. Fleet Call is closely monitoring the utilization of the network and can add bandwidth to the frame relay backbone link as required. The WAN carriers also offer specialized reporting to provide a detailed view of circuit utilization.

Issues

Fleet Call employed the services of HP Network Consultants in the design and initial rollout of the network. However, as with any network implementation, unanticipated issues surfaced. When implementing a new network, remember the following:

1. When planning a frame relay network, make sure the network equipment has been completely certified by the circuit provider. This will aid in troubleshooting and support.

2. When purchasing newly released equipment, plan the rollout of the equipment after suitable testing has been performed.

3. Never underestimate the time and effort required to bring up a new WAN link.

4. In remote-site installations, perform site walkthroughs to determine any unforeseen obstacles that may delay or impair installation of the network components.

5. Plan ahead and purchase early. You never know when network equipment lead times suddenly will be longer than needed.

6. Keep everyone informed. Develop a mutually agreed-upon schedule of implementation tasks and responsibilities. Set expectations correctly.

Index

!

10Base-FL
    See fiber-optic connection
10Base-T... 1-10
10Base2... 1-10
10Base5... 1-10
200 series specifics... 1-3, 1-10, 1-21–1-23
400 series specifics... 1-3, 1-10, 1-21–1-23
600 series ... 1-3, 1-9, 1-23, 1-25–1-26
650 specifics... 1-3, 1-9–1-10, 1-23–1-28
    See also HP J2430A Router 650
802.x
    See IEEE 802.x

A

AARP Probe... 2-113
AARP: AppleTalk Address Resolution Protocol... 1-8, 2-113
Accept NetBIOS Broadcasts configuration... 2-104–2-105
adapter... 1-11, 1-13
    See also transceivers
Address Mapping Table (AMP) for AppleTalk... 2-113
address resolution... 1-8, 2-69, 2-72, 2-76, 2-90, 2-94, 2-118
Address Resolution Protocol
    See ARP
address resolution protocols... 1-6, 1-8, 1-32
address table... 1-16–1-17, 1-20, 1-29, 2-4, 2-85, 2-93, 2-124
addressing... 1-4–1-6, 1-8, 2-3–2-4, 2-57–2-64, 2-80, 2-91, 2-96–2-100, 2-110, 2-112–2-113, 2-116–2-119
adjacency... 1-32, 2-120–2-121, 2-123
adjacent host route... 2-69, 2-76
advertisements... 2-66–2-67, 2-74, 2-79
AEP: AppleTalk Echo Protocol... 2-113–2-114
aging routes... 2-65, 2-101
all nets broadcast... 2-104–2-105
AMP: Address Mapping Table (AppleTalk)... 2-113
ANSI terminal... 1-17
Apollo Domain routing... 1-7
AppleTalk
    addressing scheme... 1-6, 2-110
    node address... 2-110, 2-113–2-114
    routing service... 1-3, 1-6–1-8, 1-19, 2-107–2-108, 2-110–2-114
    topology example... 2-107–2-109
    zone... 2-111
AppleTalk Address Resolution Protocol... 1-8, 2-113
AppleTalk Echo Protocol... 2-113–2-114
application filtering for IP... 2-81
architecture... 1-21–1-34
area number... 2-117–2-118
areas
    DECnet... 1-7, 2-116–2-117, 2-121–2-123
    OSPF... 1-6, 1-32, 2-66
ARP: Address Resolution Protocol... 1-6, 1-8, 1-32, 2-69, 2-72, 2-76, 2-90, 2-94
ARPA/Berkeley Services... 1-6
ARPAnet
    See Internet
Atping command in NCL... 2-114
AUI ports... 1-10, 1-24
autoconfiguration
    See SmartBoot
autodetection... 1-14, 1-16
autodiscovery... 1-20
automatic backup
    See backup link
autonomous system (in IP)... 2-66–2-67

B

backbone router for AppleTalk... 2-108
backup link... 1-4, 1-13, 1-29, 1-31, 2-3, 2-72
backup power... 1-24
battery-backed static RAM... 1-22
bindery (SAP)... 2-103
bindery service... 2-96
block diagram
    hardware architecture... 1-22, 1-27–1-28
BNC ports... 1-10, 1-24
booting... 1-15, 2-85–2-86
Bootp... 1-15, 1-19, 2-85–2-89
branch routing... 1-9, 1-13–1-16, 1-18–1-19, 2-85–2-86
bridge... 2-55
bridging performance... 1-26
bridging service... 1-3–1-5, 1-8, 1-13–1-15, 1-17–1-19, 1-29–1-32, 2-3–2-55, 2-90–2-91, 2-93, 2-107, 2-115
bridging tables
    See address table
broadcast... 2-103, 2-105
    NetBIOS broadcast... 2-104–2-105
Broadcast Routing Timer configuration... 2-123

buffer... 1-23, 1-27–1-28
bus... 1-22–1-23, 1-25–1-26, 1-28

C

cables
  RS-232... 1-11–1-12
  RS-422/449... 1-11–1-12
  V.24/V.28... 1-11–1-12
  V.35... 1-11–1-12
  V.36... 1-11–1-12
  X.21... 1-11–1-12
cache... 1-25–1-26, 1-28, 1-32, 2-120
  See memory
cell-interleaving packet bus... 1-25
channel service unit
  See DSU/CSU
chassis... 1-23–1-24
circuit... 1-12, 1-29, 1-31
circuit group... 1-29, 1-31–1-33, 2-68, 2-73
  See also network interface
circuit group manager... 1-30–1-33
circuit-switching... 1-11, 1-13
classes of IP addresses... 2-58–2-60, 2-62
clocking... 1-12, 1-16, 1-22–1-23, 1-28
CMOT... 2-56
communications standards... 2-94, 2-124
compression... 1-11, 1-16, 1-25, 1-30
Computer Interconnect bus (CI)... 2-115
conditional static route... 2-69, 2-73, 2-75
configuration
  Accept NetBIOS Broadcasts... 2-104–2-105
  automatic... 1-14–1-16, 1-18–1-19, 2-85–2-86
  Broadcast Routing Timer... 2-123
  DECnet costs... 2-119
  Default Zone Name... 2-112
  Deliver NetBIOS B'casts to Net... 2-104
  distribution... 1-16
  downloading or uploading... 1-15–1-17, 1-19, 2-85–2-86
  file... 1-15–1-16
  general services... 1-14, 1-17–1-18
  hardware... 1-14
  IP host-only mode... 2-90–2-92
  Local LAN Address... 2-119
  network number range... 2-112
  OSPF... 2-67
  remote... 2-85
  RIP... 2-66
  Router Priority... 2-120
  storage... 1-15–1-16, 1-19, 1-22, 1-26, 1-28, 2-85–2-86
  See also topology
  utilities... 1-17–1-19
  Zone Name list... 2-111–2-112
Configuration Editor... 1-18, 2-112, 2-119, 2-123
console... 1-9, 1-15–1-20, 2-85, 2-93, 2-124
  See also RS-232 console port
controller chip... 1-22–1-23, 1-25, 1-27–1-28
convergence... 2-66
coprocessor
  See processor chip
cost... 1-32–1-33, 2-65–2-67, 2-69–2-71, 2-74, 2-102, 2-119, 2-122
count to infinity... 2-66
CPU
  See processor chip
CSMA/CD... 2-107
CSU
  See DSU/CSU
CTERM... 2-116

D

DAP: Data Access Protocol (DECnet)... 2-116
DAS: dual attachment station... 1-11, 1-23, 1-27
Data Access Protocol (DAP) for DECnet... 2-116
data communications specifications... 2-94, 2-124
Data Link Accelerator (DLA)... 1-25, 1-27
data service unit
  See DSU/CSU
data-link service... 1-14, 1-16, 1-30–1-33
Datagram Delivery Protocol
  See DDP
daughter board... 1-21
DCE... 1-16
DDD: direct distance dialing
  See PSTN
DDN: U.S. Defense Data Network... 1-8, 1-12, 2-94
DDP: Datagram Delivery Protocol... 1-8, 1-30, 2-113–2-114
DECnet
  addressing scheme... 2-116–2-119
  device management... 2-124
  network architectures... 2-115, 2-124
  OSI architecture... 2-115
  routing service... 1-3, 1-7–1-8, 2-115–2-124
  SNMP management... 2-124
  topology example... 2-117
DECnet Routing Protocol
  See DRP
default route... 2-69–2-71, 2-75
Default Route Listen... 2-70–2-72
Default Route Supply... 2-70–2-72
Default Zone Name configuration... 2-112
Deliver NetBIOS B'casts to Net configuration... 2-104
designated router... 2-66–2-67, 2-120–2-121

device management... 1-17, 2-85, 2-90, 2-93–2-94, 2-124
DFS: Distributed File System (DECnet)... 2-116
DHCP: Dynamic Host Configuration Protocol... 2-86–2-89
diagnostic services (NetWare)... 2-96
diagnostics... 1-17
dialup lines... 1-11, 1-13
Digital Network Architecture (DECnet)... 2-115, 2-124
direct route... 1-32–1-33
directory service... 2-96, 2-114, 2-116
diskless nodes... 2-115
Distributed File System (DFS) for DECnet... 2-116
distributed file system (NetWare)... 2-96
Distributed Naming Service (DNS) for DECnet... 2-116
Distributed Queuing Service for DECnet... 2-116
DLA
See Data Link Accelerator
DNA: Digital Network Architecture (DECnet)... 2-115, 2-124
DNS: Distributed Naming Service (DECnet)... 2-116
dotted decimal notation... 2-60–2-61
download
See file transfer
driver... 1-30–1-33
DRP: DECnet Routing Protocol... 1-8, 1-30, 2-115
DSU/CSU: data service unit/channel service unit... 1-11–1-12
DTE... 1-16
Dual Attachment Station
See DAS
dual homing... 1-11
dual switch module for FDDI
See optical bypass switch
dynamic link configuration... 1-13–1-14, 1-16
dynamic routing... 2-65, 2-69, 2-101

E

EASE sampling... 1-25
EASE: Embedded Advanced Sampling Environment... 1-20
Echo Protocol... 1-7
EEPROM... 1-16, 1-22, 1-26, 1-28
EGP: Exterior Gateway Protocol... 1-6, 1-8, 1-32, 2-65, 2-67, 2-74, 2-94
electronic mail... 2-56, 2-96
encapsulation for adjacent host routes... 2-76, 2-78
encapsulation for Novell... 2-99
engine... 1-24–1-26, 1-28
EPROM... 1-22, 1-28
Error Protocol... 1-7
Ethernet... 1-3–1-7, 1-10, 1-23–1-24, 1-27, 2-3, 2-77, 2-93, 2-107–2-108
EtherTalk... 2-107
event log... 1-16–1-17, 2-93, 2-124
export route filter... 1-32, 2-79–2-80, 2-94
Exterior Gateway Protocol
See EGP

F

factory default configuration... 1-15, 1-18
fan... 1-9, 1-21, 1-24
fast-access static RAM... 1-22, 1-27–1-28
FCS
See frame check sequence
FDDI... 1-3–1-4, 1-11, 1-23, 1-27, 2-3
features... 1-3–1-20
fiber-optic connection... 1-23–1-24, 1-27
FDDI... 1-11
IEEE 802.3 Type 10Base-FL or FOIRL... 1-10, 1-24
fiber-optic wavelength... 1-11
file server... 1-15, 2-85–2-86, 2-95–2-96, 2-102
file transfer... 1-15–1-17, 2-56, 2-81, 2-85, 2-94, 2-116
filtering... 1-4, 1-8, 1-25, 1-30, 1-32–1-33, 2-3, 2-79–2-84, 2-94, 2-103
Fixed Shroud Duplex MICs
See FSD
flash EEPROM... 1-16, 1-22, 1-26, 1-28
flow control... 1-31
FOIRL
See fiber-optic connection
FOIRL: 802.3 fiber-optic connection... 1-24
forwarding table
See routing table
frame check sequence (FCS)... 1-23, 1-27
frame formatting... 1-23, 1-27
frame relay... 1-3, 1-11
FSD: Fixed Shroud Duplex... 1-11
FTP: File Transfer Protocol... 2-56, 2-81

G

garage... 1-24
gateway... 2-67
general service response
See GSR
Graybook... 1-7
GSI: Government Systems, Inc.... 2-57, 2-59, 2-94
GSR: general service response (IPX)... 2-103

H

half router for AppleTalk... 2-108
hardware architecture... 1-9, 1-21–1-28
HDLC... 1-16, 2-123
Hello messages for DECnet... 2-121
help... 1-18
hierarchical routing (DECnet)... 2-116–2-118, 2-121–2-123
hop count... 1-5, 2-65–2-66, 2-100, 2-102–2-104
host number... 2-96–2-98, 2-100, 2-102–2-103
host-only bridging for IP... 2-90–2-93
hot swap... 1-9, 1-24–1-25
HP 28674B Remote Bridge RB... 1-13–1-14, 1-16, 2-63
HP Ease
    See EASE
HP J2430A Router 650... 1-15, 1-19
    hardware... 1-3, 1-9–1-10, 1-23–1-28
HP J2540A Router PR... 1-15, 1-19
HP OpenView Interconnect Manager... 1-20, 2-63, 2-85, 2-90
HP Probe... 1-6, 1-8, 2-69, 2-72, 2-90, 2-94, 2-125–2-132
hub, a network device... 2-55, 2-91

I

i960... 1-25, 1-27–1-28
IBM computers... 2-56
ICMP: Internet Control Message Protocol... 1-6, 2-94
IDP: Internet Datagram Protocol... 1-7, 1-30
IEEE 802.1
    See spanning tree
IEEE 802.3... 1-3–1-5, 1-10, 1-23–1-24, 1-27, 2-3, 2-77, 2-93
IEEE 802.5... 1-4, 1-7, 1-10, 1-23, 1-27, 2-3, 2-93
    See also token ring
IGP: interior gateway protocol... 1-5–1-6, 1-8, 2-66–2-67
IGRP on Cisco routers... 2-71
import route filter... 1-32, 2-79–2-80, 2-94
Instant On... 1-14–1-16, 2-85
inter-area routing... 2-121–2-123
Interconnecting With TCP/IP reference... 2-57
interface
    See network interface
interface card module... 1-24–1-25
interior gateway protocol
    See IGP
Internal LAN ID (HEX)... 2-47
Internet... 1-5, 2-58–2-59, 2-67, 2-93–2-94
Internet Activities Board... 2-94
Internet Control Message Protocol
    See ICMP
Internet Datagram Protocol
    See IDP
Internet Packet Exchange
    See IPX
Internet Protocol
    See IP
internetwork... 1-3, 1-5, 2-56, 2-58, 2-63, 2-93, 2-96–2-98, 2-101, 2-104–2-105, 2-107, 2-110–2-111, 2-114, 2-116–2-117, 2-119, 2-121
Internetwork Packet Exchange
    See IPX
intra-area routing... 2-121–2-123
invalid route... 2-65, 2-101
IP
    address... 1-13, 2-57–2-64, 2-68–2-69, 2-81, 2-83, 2-90–2-91
    addressing scheme... 1-5, 2-57–2-64, 2-67
    applications... 2-56, 2-81
    autonomous system... 2-66–2-67
    device management... 2-85, 2-90, 2-93–2-94
    end node... 1-5, 2-92–2-93
    filter... 2-79–2-84, 2-94
    host... 2-58–2-62, 2-64, 2-90, 2-92–2-93
    OSPF... 1-5–1-6, 1-8, 1-32, 2-66–2-67, 2-74
    performance... 1-26
    routing protocols... 1-5, 1-8, 2-66–2-67, 2-69–2-71, 2-74–2-75, 2-79–2-80
    routing service... 1-3, 1-5–1-6, 1-8, 1-15, 1-19, 1-32, 2-55–2-94
    routing table... 1-32 , 2-64–2-65, 2-67, 2-74–2-75, 2-80, 2-85
    SNMP management... 2-56, 2-85, 2-90, 2-93–2-94
    subnet mask... 2-61–2-64, 2-67
    subnet number... 2-58–2-59, 2-64
    topology example... 2-63
    virtual IP host on non-IP networks... 2-90–2-93
IPX
    addressing scheme... 2-96–2-100, 2-119
    NetWare services... 2-95–2-96, 2-102–2-104
    Novell encapsulation... 2-99
    packet... 2-96–2-97, 2-100
    performance... 1-26
    routing service... 1-3, 1-6, 1-8, 1-15, 1-19, 1-33, 2-92, 2-95–2-106
    routing table... 1-33, 2-100–2-102
    SAP... 1-33
    topology example... 2-90–2-91, 2-98–2-99
IPXWAN... 1-6, 2-95–2-96
ISDN... 1-3, 1-11, 1-13
ISO
    See OSI

IW2
See IPXWAN

K

Kerberos... 2-56

L

LAN types supported... 1-3, 1-10–1-11, 1-24, 2-95, 2-107
LAT: Local Area Transport... 2-115
latency... 1-11, 1-13, 1-16, 1-25
learning bridge... 1-3–1-4, 2-3
learning routes... 2-65, 2-101
LEDs... 1-9, 1-24
Level 1 and Level 2 routing (DECnet)... 1-7, 2-121–2-123
levels of SAP filters... 2-103
line... 1-29
link... 1-29
  costs... 2-119, 2-122
  negotiation... 1-14, 1-16
  See also network interface
  See also point-to-point WAN link
  speed... 1-16
  See also X.25
link-state advertisement (LSA)... 2-66–2-67, 2-74
link-state database... 1-32
link-terminating equipment
  See WAN link terminating equipment
LLC1 and LLC2... 1-16
load sharing
  circuits... 1-11, 1-13, 1-29, 1-31
  power supply... 1-24
Local Area Transport (LAT)... 2-115
local router for AppleTalk... 2-107
Local Zone Table for AppleTalk... 2-114
LocalTalk... 1-6–1-7
logged events
  See event log
loops, routing... 2-66

M

MAC address
  See station address
Maintenance Operations Protocol (MOP)... 2-115
management bus... 1-25–1-26
management of networks
  See network management
management of routers
  See network management or device management
manual adapter dialing... 1-11, 1-13
map of network... 2-63, 2-98–2-99, 2-109, 2-117, 2-119
media interface connector
  See MIC
memory... 1-22–1-23, 1-25–1-28
memory bus... 1-25
Message Handling Service (MHS) for NetWare... 2-96
MHS: Message Handling Service (NetWare)... 2-96
MIB... 1-17, 1-20, 2-85, 2-124
MIC: media interface connector for FDDI rings... 1-11
MIC: medium interface connector
  for token rings... 1-10
modem... 1-11–1-12, 1-15, 1-17
modules for 650... 1-24–1-25
MOP: Maintenance Operations Protocol... 2-115
motherboard... 1-21
multicast address... 1-6
multimode fiber... 1-11

N

Name Binding Protocol (AppleTalk)... 2-113–2-114
naming service
  See directory service
NBP: Name Binding Protocol (AppleTalk)... 2-113–2-114
NCL: Network Control Language Interpreter... 2-93, 2-124
  Atping command... 2-114
  Rget commands... 2-103, 2-114
  Rgetat command... 2-114
  Rgetis command... 2-103
NCP: NetWare Core Protocols... 2-95, 2-102
NetBIOS... 1-4, 2-3, 2-56, 2-104–2-105
Netware
  See IPX
NetWare Core Protocols
  See NCP
network address... 1-8, 1-29, 2-63–2-64, 2-96, 2-98, 2-100, 2-103, 2-110, 2-118
network diameter... 2-66
Network File Transfer (NFT) for DECnet... 2-116
network interface... 1-29, 1-31, 2-55, 2-64–2-65, 2-68, 2-70–2-71, 2-81, 2-90–2-91, 2-98, 2-101, 2-103–2-104, 2-120
network management... 1-20, 1-31, 2-56, 2-63, 2-85, 2-90–2-92, 2-94, 2-105
network map... 2-63, 2-98–2-99, 2-109, 2-117, 2-119
network number... 2-58–2-62, 2-79–2-80, 2-96–2-98, 2-100–2-103, 2-110, 2-112, 2-114
next hop... 2-65, 2-69–2-71, 2-73, 2-100–2-102
NFS: Network File System... 2-56
NFT: Network File Transfer (DECnet)... 2-116

node... 1-8, 1-17, 2-55, 2-58–2-59, 2-64, 2-68, 2-95, 2-97–2-98, 2-107–2-108, 2-110–2-114, 2-116–2-118, 2-120–2-121
node bypassing in FDDI... 1-11
node identifier
See node number
node number... 2-110, 2-113–2-114, 2-117–2-120
Normal ARP... 2-72
Novell
See IPX
NS Services... 1-6
NTP: network time protocol... 2-56

O

online replacement of hardware
See hot swap
operating system... 1-14, 1-16, 1-22, 1-26, 1-28, 2-97
downloading or uploading... 1-14, 1-16–1-17, 1-22, 1-28, 2-85–2-86, 2-115
operations on router... 2-93, 2-114, 2-124
optical bypass switch... 1-11
OS
See operating system
OSI... 2-115
OSI reference model... 1-29, 1-31, 2-113
OSPF: Open Shortest Path First... 1-5–1-6, 1-8, 1-32, 2-65–2-67, 2-74, 2-94

P

packet
buffer... 1-22–1-23, 1-25, 1-27–1-28
bus... 1-25
filter... 1-30, 2-81, 2-83
switching... 1-11–1-12, 1-25, 2-63, 2-108
Packet Exchange Protocol
See PEP
packet headers and addresses... 1-29, 1-31, 2-84
pass-through of synchronous traffic... 1-13, 2-119
password
console... 1-17
OSPF... 1-6, 2-67
path cost... 2-119, 2-122–2-123
PC used as a console device... 1-16–1-17
PCA: printed circuit assembly... 1-9, 1-21, 1-24
PCMCIA flash card... 1-26, 1-28
PDN: Public Data Network... 1-8, 2-94
PEP: Packet Exchange Protocol... 2-95, 2-104
performance... 1-25–1-26, 1-30–1-31
permanent virtual circuit... 1-12
ping: Packet InterNet Groper... 2-85, 2-90, 2-94, 2-114
planning... 2-117, 2-119
point-to-point WAN link... 1-11–1-13, 1-19, 2-72, 2-123
poison reverse (in RIP)... 2-66
port
controller... 1-22–1-23, 1-25, 1-27–1-28
filter... 2-81, 2-94
POTS
See PSTN
power
supply... 1-9, 1-21, 1-24
PPP: Point-to-Point Protocol... 1-13
preference... 1-32, 2-74–2-75
print server... 2-96, 2-102, 2-116
printed circuit assembly
See PCA
printer queue... 2-96, 2-116
prioritization... 1-30
prioritization of WAN traffic... 1-11, 1-13
priority of router for DECnet
See designated router
Privacy Enhanced Mail... 2-56
Probe: AARP Probe for AppleTalk... 2-113
Probe: HP Probe for IP... 1-6, 1-8, 2-69, 2-72, 2-90, 2-94, 2-125–2-132
processor chip... 1-22–1-25, 1-27–1-28
PROM: programmable read-only memory... 1-22, 1-27
protocols
Address Resolution Protocol... 1-6, 1-8, 2-69, 2-72, 2-76, 2-90
AppleTalk Address Resolution Protocol... 1-8, 2-113
AppleTalk Echo Protocol... 2-113–2-114
Bootp... 1-15, 1-19, 2-85–2-89
Data Access Protocol for DECnet... 2-116
Datagram Delivery Protocol (DDP)... 1-8, 1-30, 2-113–2-114
DDCMP for DEC routers... 2-123
DECnet Routing Protocol... 1-8, 2-115
DHCP... 2-86–2-89
EGP... 1-6, 1-8, 2-65, 2-67, 2-74
HDLC... 2-123
HP-proprietary point-to-point... 1-13
Internet Datagram Protocol (IDP)... 1-8, 1-30
IPX... 1-8, 2-95, 2-104
IPXWAN... 1-6, 2-95–2-96
LAT... 2-115
MOP... 2-115
Name Binding Protocol for AppleTalk... 2-113–2-114
NetBIOS... 2-56, 2-104–2-105
NetWare Core Protocols (NCP)... 2-95, 2-102
OSPF... 1-5–1-6, 1-8, 2-65, 2-67, 2-74, 2-94
Packet Exchange Protocol (PEP)... 2-95, 2-104
Point-to-Point Protocol (PPP)... 1-13

RIP... 1-5, 1-7–1-8, 2-65–2-67, 2-69–2-70, 2-74, 2-79–2-80, 2-94
See also routing protocols and routing services
Routing Table Maintenance Protocol... 1-8, 2-112–2-114
SNMP... 1-20
spanning tree... 1-3–1-4, 2-3
spanning tree protocol... 1-29
SPX... 2-95
System Communication Services... 2-115
TFTP... 1-15–1-16, 2-81, 2-85
Zmodem... 1-16
Zone Information Protocol for AppleTalk... 2-113–2-114
PSTN: Public Switched Telephone Network... 1-11
PVC: permanent virtual circuit... 1-12

Q

QMS: Queue Management Services (NetWare)... 2-96
quality of service... 1-16
Queue Management Services (QMS) for NetWare... 2-96
Quick Configuration... 1-18, 2-112, 2-119
Quick Remote... 1-15, 1-19, 2-85

R

rack mounting... 1-9
RAM: random-access memory... 1-22–1-23, 1-26–1-28
random-access memory
See RAM
reachability of nodes... 2-93–2-94
read-only memory
See ROM
Record Management Service for DECnet... 2-116
redirector... 1-29–1-33
redundant power supply... 1-9, 1-24
remote terminal
See Telnet
See also CTERM (DECnet)
repeater... 2-55, 2-91
Reset button... 1-23–1-24
RFC
1010... 2-94
1058... 2-94
1155... 2-94
1156... 1-20, 2-85, 2-94
1157... 2-94
1247... 2-94
783... 2-94
792... 2-94
826... 2-94
854... 2-94
868... 2-94
877... 1-8, 1-12, 2-94
Rgetat command... 2-114
Rgetis command... 2-103
ring, FDDI
See FDDI
ring, token
See token ring
RIP Listen... 2-70–2-72
RIP Supply... 2-70–2-72
RIP: Routing Information Protocol... 1-5, 1-7–1-8, 1-31–1-33, 2-65–2-67, 2-69–2-70, 2-74, 2-79–2-80, 2-94, 2-101–2-102
RISC... 1-25, 1-27–1-28
ROM: read-only memory... 1-22, 1-27
routable protocols and suites
See routing services
route filter... 2-79–2-80
Route Learned
IP routing table... 2-65
IPX routing table... 2-101
Route Type
IP routing table... 2-65
IPX routing table... 2-101
router interface
See network interface
Router Priority
See designated router
routing decisions... 1-31–1-33, 2-64–2-67
Routing Information Protocol
See RIP
routing metric for DECnet... 2-122
routing performance... 1-26
routing pool... 2-74, 2-79–2-80
routing protocols
DRP (DECnet)... 1-8, 1-30, 2-115
EGP... 1-6, 1-8, 1-32, 2-65, 2-67, 2-74, 2-94
IGRP on Cisco routers... 2-71
OSPF... 1-5–1-6, 1-8, 1-32, 2-65–2-67, 2-74, 2-94
RIP... 1-5, 1-7–1-8, 1-31–1-33, 2-65–2-67, 2-69–2-70, 2-74, 2-79–2-80, 2-94, 2-101–2-102
See also routing services
RTMP for AppleTalk... 1-8, 1-31, 2-113–2-114
routing services... 1-3–1-8, 1-15, 1-17–1-19, 1-29–1-32, 2-55–2-94, 2-124
AppleTalk Phase 2... 1-6–1-8, 1-19, 2-107–2-108, 2-110–2-114
DECnet... 1-7–1-8, 2-115–2-124
Internet Protocol (IP)... 1-5–1-6, 1-8, 1-15, 1-19, 1-32, 2-55–2-94

Novell IPX... 1-6, 1-8, 1-15, 1-19, 1-33, 2-92, 2-95–2-106
XNS... 1-7–1-8
routing table... 1-16–1-17, 1-20, 1-29, 1-32–1-33, 2-64–2-65, 2-67, 2-74–2-75, 2-80, 2-85, 2-93, 2-100–2-102, 2-105, 2-114, 2-120–2-123
RS-232 console port... 1-9, 1-15–1-24, 1-28
RS-232 WAN interface... 1-11–1-12
RS-422 WAN interface... 1-11–1-12
RS-449 WAN interface... 1-11–1-12
RTMP: Routing Table Maintenance Protocol... 1-8, 1-31, 2-112–2-114

S

sampling of traffic... 1-20
SAP: Service Advertising Protocol... 1-33, 2-102–2-103
filter... 1-33, 2-103
table for IPX... 2-103
SAS: service advertising socket... 2-103
SAS: single attachment station... 1-11
SCA: System Communications Architecture... 2-115
security... 1-5–1-6, 1-8, 1-17, 2-56, 2-67, 2-72, 2-105
seed router... 2-110–2-112
segment... 2-55
self-test... 1-22, 1-28
Sequenced Packet Exchange
See SPX
serial port
See console
serial ports and links
See WAN
series
See 200, 400, or 600
server... 2-95–2-96, 2-102–2-105, 2-114
Service Advertising Protocol
See SAP
service advertising socket (SAS)... 2-103
Simple Network Management Protocol
See SNMP
Single Attachment Station
See SAS
site survey... 2-119
slow convergence... 2-66
SmartBoot... 1-14–1-16, 1-18–1-19, 2-85–2-86
SMDS... 1-3, 1-11
SMTP: Simple Mail Transfer Protocol... 2-56
SNAP... 2-78
SNMP: Simple Network Management Protocol... 1-20, 1-31, 2-56, 2-85, 2-90–2-94
socket number (IPX)... 2-96–2-97, 2-102–2-103
software
architecture... 1-16, 1-29–1-30, 1-32–1-33
distribution... 1-16
See also operating system or configuration
SONIC... 1-23, 1-27
source routing... 1-5, 1-29, 2-93, 2-115
source-routing bridge... 1-4–1-5, 2-3, 2-93
spanning tree... 1-3–1-4, 1-29, 1-31, 2-3
specifications... 2-94, 2-124
split horizon (in RIP)... 2-66
SPX: Sequenced Packet Exchange... 1-6, 2-95
SRAM: static random-access memory... 1-22, 1-27–1-28, 1-34
SRT: source-routing/transparent bridge... 1-4, 2-3
static RAM
See SRAM
static route... 1-8, 1-32–1-33, 2-65, 2-67, 2-69–2-78, 2-94, 2-101–2-102, 2-105
NetBIOS static route... 2-104–2-105
station address... 1-22, 1-27, 1-29, 2-4, 2-69, 2-97–2-98, 2-100, 2-113, 2-118–2-119
statistics... 1-16–1-17, 1-20
status indicators
See LEDs
STP: IEEE Spanning Tree Protocol
See spanning tree
subnet mask... 2-61–2-64, 2-67, 2-69, 2-80, 2-90–2-91
subnet number... 2-59–2-61
subnetwork... 1-5, 2-55, 2-59, 2-63–2-65, 2-67–2-68, 2-90, 2-98
SVC: switched virtual circuit... 1-12
Switched 56... 1-11
synchronous pass-through... 1-13, 2-119
synchronous ports
See WAN ports
system bus
See global bus
System Communication Services... 2-115
System Communications Architecture (SCA)... 2-115

T

T connector for BNC port... 1-10
T1... 1-11
TCP/IP protocol suite
See IP
TCP/UDP port filters... 2-81–2-84, 2-94
TCP: Transmission Control Protocol... 1-6
technical data... 2-124
Telnet... 1-15, 1-17, 1-19–1-20, 2-56, 2-81, 2-85, 2-90, 2-94
terminal adapter... 1-11, 1-13
terminology and basic concepts... 1-29
test reachability of nodes... 2-85, 2-94, 2-114

TFTP: Trivial File Transfer Protocol... 1-15–1-17, 2-56, 2-81, 2-85, 2-90, 2-94
thick coaxial LAN connection... 1-10
thin coaxial LAN connection... 1-10
ThinLAN... 1-10
throughput
See performance
time protocol... 2-56, 2-85, 2-90, 2-94
timer
See clocking
token ring... 1-3–1-7, 1-10, 1-23, 1-27, 2-3, 2-93, 2-107
TokenTalk... 2-107
topology... 2-64, 2-68, 2-75, 2-101
topology example
AppleTalk... 2-107–2-109
DECnet... 2-117
IP internetwork... 2-63
IPX... 2-98–2-99
IPX with virtual IP host configuration... 2-91
traffic control... 1-5, 1-8
traffic filters
See packet filter
traffic monitoring... 1-20
Transaction Tracking System (TTS) for NetWare... 2-96
transceivers... 1-9–1-10, 1-23–1-24, 1-28
translational bridging... 1-4, 2-3
Transmission Control Protocol
See TCP
transmit queue... 1-16
Transpac... 2-63
transparent bridge
See learning bridge
Trivial File Transfer Protocol
See TFTP
trunk coupling unit (TCU)... 1-10
TTS: Transaction Tracking System (NetWare)... 2-96
twisted-pair connection... 1-10, 1-24

U

UDP: User Datagram Protocol... 1-6, 2-81, 2-94
update operating system... 1-16–1-17, 1-22, 1-26, 1-28
upgrade memory... 1-26, 1-28
User Datagram Protocol
See UDP

V

V.24/V.28
See RS-232
V.25 bis adapters... 1-11, 1-13
V.35 WAN interface... 1-11–1-12
V.36
See RS-422/449
VAX... 2-115
Videotex... 2-116
virtual circuit... 1-12
virtual IP host on non-IP networks... 2-90–2-93
virtual terminal (DECnet)... 2-116
VMS... 2-116
voltage, auto-adjusting... 1-9, 1-21, 1-24
VT100... 1-17

W

WAN
cables... 1-11–1-12, 1-14
Dynamic Link Configuration... 1-14, 1-16
link speeds... 1-11–1-12, 1-16
link-terminating equipment... 1-11–1-12
load sharing... 1-11, 1-13
multiple links... 1-11, 1-13
point-to-point link... 1-11–1-13, 1-19, 2-72, 2-123
ports... 1-11–1-14, 1-23–1-24, 1-27, 2-97, 2-100, 2-107
protocols... 1-13
WAN ports... 1-25
warranty... 1-9
worksheet for planning... 2-119

X

X Window System... 2-56
X.21 WAN connection... 1-11–1-12
X.25... 1-3, 1-11–1-12, 2-63, 2-108
Xerox Network Systems Internet Transport Protocols suite
See XNS
XNS
routing service... 1-3, 1-7–1-8, 2-119

Z

ZIP: Zone Information Protocol (AppleTalk)... 2-113–2-114
ZIT
See Zone Information Table
Zmodem... 1-16
zone... 2-111–2-112, 2-114
Zone Information Protocol (ZIP) for AppleTalk... 2-113–2-114
Zone Information Table (AppleTalk)... 2-114
Zone Name list configuration... 2-111–2-112