Internet Protocol

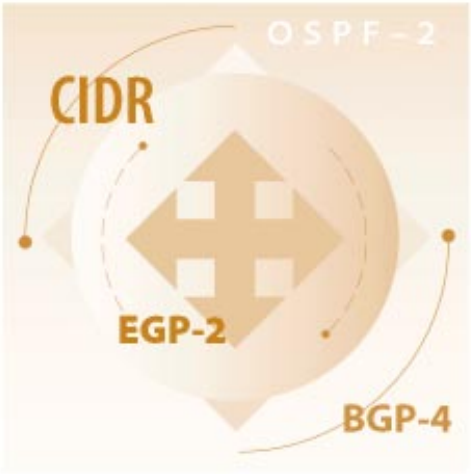

Maximizes Network Performance and Availability

Maximizes Network Connectivity

Enhances Network Security

Internet Protocol (IP), the network protocol used by corporations, governments, and the Internet worldwide, supports many personal, technical, and business applications, such as electronic funds transfers, medical records processing, image transfers, and electronic mail.

IP is the network layer protocol of the Internet Protocol (IP) suite that specifies the data format, performs routing functions, and provides an unreliable, connectionless data delivery service on a "best effort basis." IP defines the internet datagram as the unit of information passed across the internetwork.

Bay Networks IP implementation for its routers provides all standard IP functions and supports major Internet Engineering Task Force (IETF) Request For Comments (RFCs) for protocols including IP, TCP, UDP, ICMP, RIP, CIDR, OSPF, EGP-2, BGP-3, BGP-4, BOOTP Gateway, Router Discovery, ARP, RARP, SNMP, Telnet, TFTP, FTP, and RIPSO. This support ensures connectivity and interoperability, allows the internetwork to be managed effectively, and enhances network security. Bay Networks enhances IP with features that include IP Routing Policies, IP Multicast Routing support, Static Routes, Adjacent Hosts, Circuitless IP Interface Support, Configurable RIP Timers, Configurable Triggered Updates, Source Route Bridge End Station Support, NetBIOS over IP Support, Ping MIB, Crypto Resynchronization Support, and Blacker Front-End Support.

Bay Networks IP support of Ethernet, Token Ring, FDDI, ATM, and serial interfaces maximizes internetwork connectivity. The router's serial interfaces operate at rates up to 52 Mbps and support WAN links such as Frame Relay, SMDS, PPP, ATM DXI, and X.25. Dial Back-up, Dial-on-Demand, and Bandwidth-on-Demand using switched services are also supported. Multiline Circuits, Uniform Traffic Filters, Traffic Prioritization, and Data Compression optimize bandwidth and maximize IP traffic control. IP is easily configured on all Bay Networks routers via Bay Networks Optivity/Internetwork™, which combines three applications — Site Manager, RouterMan, and PathMan — to form a seamless integrated router management package for Bay Networks routers.

All Bay Networks multiprotocol router/bridges support IP. The IP implementation complements other supported network and bridging protocols including OSI, DECnet Phase IV, Novell IPX, Banyan VINES®, AppleTalk Phase 2®, XNS, APPN, Data Link Switching (DLSw), Native Mode LAN (NML), Translation Bridge, Transparent Bridge, and Source Route Bridge.

Benefits

Maximizes Network Performance and Availability

Bay Networks IP implementation maximizes network throughput and accessibility by optimizing network efficiency. By supporting Classless Inter-Domain Routing (CIDR), Bay Networks routers guard against the exhaustion of Class B network address space and ensures manageability in large networks. OSPF supports small and infrequent routing updates and fast network convergence. Using TCP, BGP-3, and BGP-4 provide reliable, loop-free routing with minimal network overhead. BGP-4 support further reduces network overhead with support for CIDR. Network performance is enhanced by Configurable RIP Interface Costs, which allows user-defined "preferred" route selection. Additionally, the router's Circuitless IP Interface feature ensures router reachability if IP interfaces are disabled. By allowing a WAN circuit to be composed of multiple data paths, Multiline Circuits increases available bandwidth and provides link redundancy. IP multicast improves bandwidth efficiency by sending packets only to specified end stations.

Maximizes Network Connectivity

Comprehensive network connectivity is ensured by the Bay Networks router's support of all major IP protocols, including OSPF, RIP, EGP-2, BGP-3, and BGP-4. Additionally, extensive IP address resolution services are provided through Address Resolution Protocol (ARP), Reverse ARP (RARP), and BOOTP Gateway support, while the Adjacent Hosts feature enhances connectivity by allowing routes to be predefined for networks and hosts not supporting ARP services.

Furthermore, the router's Source Route Bridge End Station routes IP traffic from a source route bridge environment to end stations on remote LANs over a multiprotocol backbone. IP is supported on the full range of LAN and WAN interfaces for Bay Networks routers.

Enhances Network Security

Bay Networks routers support the secure transmission of datagrams using RIPSO security labels, as well as Blacker Front-End and KG84A cryptographic devices. Additionally, the Uniform Traffic Filters, Static Routes, and IP Routing Policies features provide complete control of data flow.

Features

Comprehensive IP Support

The Internet Protocol suite corresponds to the OSI reference model. Bay Networks support of the Internet Protocol (IP) suite encompasses all standard medias and protocols as represented in the OSI reference model (see Figure 1). Bay Networks IP implementation conforms to all standards defined by the Internet Engineering Task Force (IETF) (see Table 1) and other accepted standards setting organizations.

The Internet Protocol suite is supported across the entire line of Bay Network's router and bridge products providing users a variety of choices to meet a broad range of networking demands.

Internet Protocol (IP) IP is a connectionless datagram delivery protocol that performs addressing, routing and control functions for transmitting and receiving datagrams over a network. As a connectionless protocol, IP does not require a predefined path associated with a logical network connection. As packets are received by the router, IP addressing information is used to determine the best "next hop" a packet should take enroute to its final destination. As a result, IP does not control data path usage. If a network device or line becomes unavailable, IP provides the mechanism needed to route datagrams around the affected area.

IP datagrams begin with a packet header. The header identifies the version of IP protocol used to create the datagram, the header length, the type of service required for the datagram, the length of the datagram, the datagrams identification number, fragmentation control information, the maximum number of hops the datagram can be transported over the internetwork, the protocol format of the data field, the source and destination addresses, and potentially IP options (see Figure 2).

Transmission Control Protocol (TCP)

TCP provides a reliable, connection-oriented, transport layer link between two hosts. Using a two-way handshaking scheme, TCP provides the mechanism for establishing, maintaining, and terminating logical connections between hosts using IP as its transport protocol. Additionally, TCP provides protocol ports to distinguish multiple programs executing on a single device by including the destination and source port number with each message. TCP performs functions such as reliable transmission of bytes streams, data flow definitions, data acknowledgments, lost or corrupt data re transmissions and multiplexing multiple connections through a single network connection. Furthermore, TCP is responsible for encapsulating information into a datagram structure.

Bay Networks IP implementation supports extension to TCP, per RFC1323 (Van Jacobson TCP), including Windows Scale, a fast retransmit/fast recovery algorithm, and round-trip time measurement. These extensions improve performance and provide reliable operations over high-speed paths.

Figure 1 **Bay Networks IP Protocol Support**

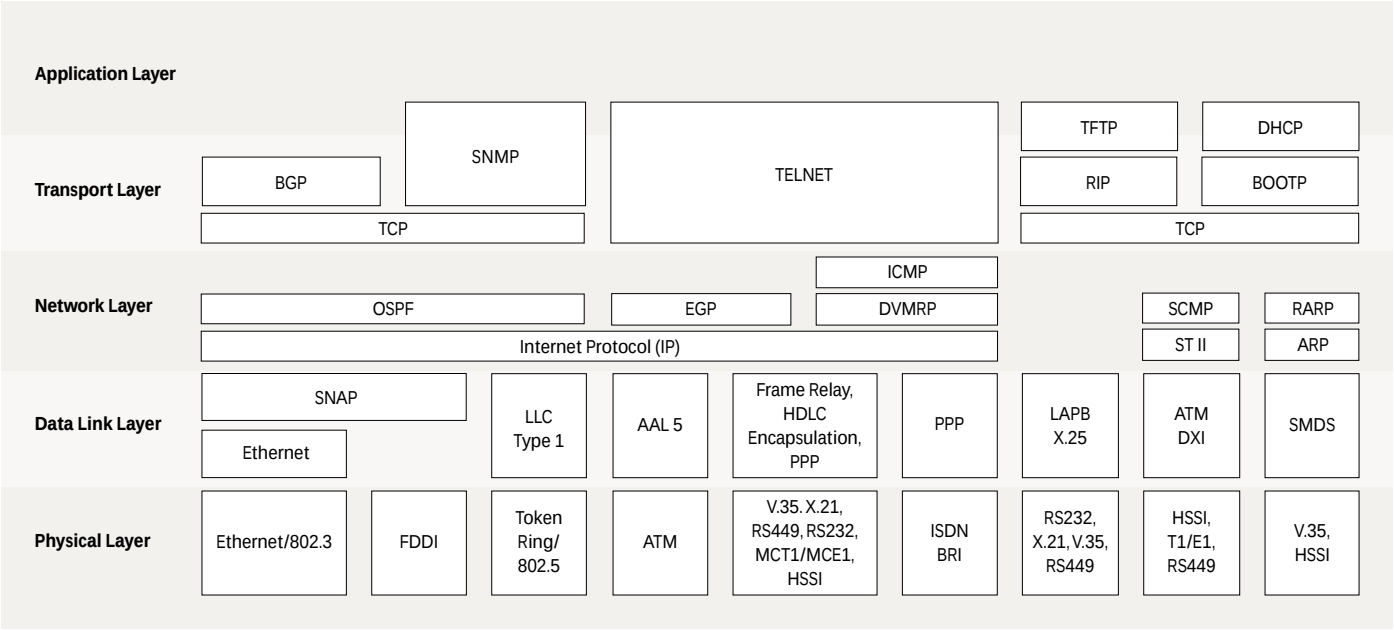

Figure 2 **IP Datagram Format**

Version	Header Length	Service Type	Total Length	Identi-fication	Flags	Fragment Offset	Time to Live	Protocol	Header Checksum	Source IP Address	Desti-nation IP Address	IP Options	Padding	Data
(4 bits)	(4 bits)	(1)	(2)	(2)	(3 bits)	(13 bits)	(1)	(1)	(2)	(4)	(4)	(variable)	(variable)	(0 to 65,500)

(# bytes)

Figure 3 | **Routing Protocol Usage**

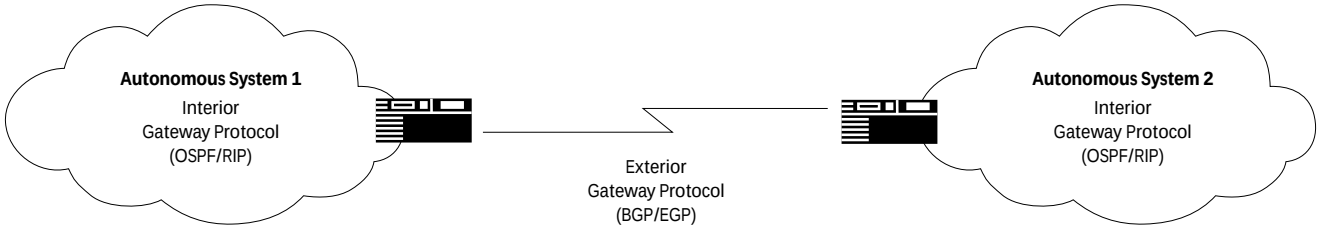

User Datagram Protocol (UDP) UDP provides a connectionless datagram delivery service between IP host applications. This protocol is used for transaction-oriented utilities such as the IP standard Simple Network Management Protocol (SNMP) and Trivial File Transfer Protocol. Like TCP, UDP encapsulates information into a datagram structure, works with IP to transport message to a destination, and provides protocol ports to distinguish between software applications executing on a single host. Unlike TCP, however UDP avoids the overhead of reliable data transfer mechanism by not protecting against datagram loss or duplication.

Routing Information Protocol (RIP) RIP is a standards-based, distance-vector, interior gateway protocol (IGP) used by routers in an autonomous system (AS) to exchange routing information (see Figure 3). Through RIP, end stations and routers are provided with the information required to dynamically choose the best paths to different networks throughout the environment. RIP implementation relies on the total number of hops between a source and destination network as the cost variable used in making best path routing decisions. The network path providing the fewest number of hops between the source and destination network is considered the path with the lowest overall cost.

The maximum allowable number of hops a packet can traverse in an IP network implementing RIP is 15 hops. By specifying a maximum number of hops, RIP avoids the occurrence of routing loops. A datagram is routed through the inter network via an algorithm that uses a routing table (RIP table) in each router. A router's RIP table contains information on all known networks in the autonomous system. The RIP table includes the total number of hops (hop count) to a destination network and the address of the "next hop" router in the direction of the destination network.

In a RIP network, each router broadcasts its entire RIP table to its neighboring router every 30 seconds. When a router receives a neighbor's RIP table, it uses the information provided to update its own routing table and then sends the updated table to its neighbors.

This procedure is repeated by each router and results in a state referred to as network convergence, in which all routers have an identical view of the internetwork topology.

Open Shortest Path First Version 2 (OSPF-2) OSPF is a second generation standards-based IGP that enables routers in an autonomous system to exchange routing information. Developed as an alternative to RIP, OSPF minimizes network convergence times within large IP internetworks.

Routers supporting OSPF exchange routing information within an autonomous system using a link state algorithm. Protocols based on link-state algorithms issue routing update messages only when a change in topology occurs. When a topology change occurs, affected routers immediately multicast information about the topology change only, instead of the entire routing table. This reduces the amount of traffic on the inter network. Additionally, since topology change information is propagated immediately, network convergence is achieved more quickly than relying on the timer based mechanism used with RIP.

With OSPF, autonomous systems can be segmented into areas — a group of contiguous networks, hosts, end stations, and routers. The use of areas reduces internetwork overhead by decreasing the amount of topology change information transmitted within the autonomous system. OSPF also provides the ability to assign user-configurable cost metrics to each router interface, allowing preferred paths to be specified for enhanced traffic control. Additionally, OSPF's variable length subnet mask feature increases network efficiency by allowing the network to be

Figure 4 | **OSPF/RIP Coexistence**

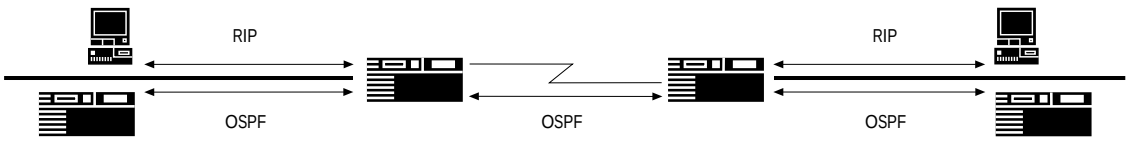

divided into subnets of varying sizes. The reachability of multiple networks within an area can be summarized through OSPF's route summarization feature. This reduces link advertisements overhead by allowing one link advertisement to be generated for all subnets in an area.

OSPF's is increasingly being adopted within existing autonomous system that previously relied on RIP's routing services. OSPF routers simultaneously support RIP for router-to-end station communications, and OSPF for router-to-router communications. Bay Networks IP implementation supports OSPF/RIP coexistence, which allows OSPF and RIP to share routing information (see Figure 4). This ensures communications within an inter network and provide a smooth migration path for introducing OSPF into existing networks.

Classless Inter-Domain Routing (CIDR) CIDR defines a strategy for IP address assignment. This strategy is an attempt to conserve the address space and to slow the explosive growth of the routing table. CIDR removes the concept of class from IP addresses.

An example of a CIDR aggregated address, referred to as a supernet address, is 192.32.0.0/16, where 192.32.0.0 represents the address prefix, and 16 is the prefix length in bits. Such an address represents destinations from 192.32.0.0 to 192.32.255.255. CIDR is supported by OSPF and BGP-4.

Border Gateway Protocol Version 4 (BGP-4) BGP-4 is an exterior gateway protocol that enables routers in different autonomous systems to exchange routing information (See Figure 3). It also provides a set of mechanisms for facilitating CIDR by providing the capability of advertising an arbitrary length IP prefix and thus eliminating the concept of network "class" within BGP. BGP uses TCP to ensure delivery of inter autonomous system information. Update messages are generated only if a topology change occurs and contain only information about the change. This reduces network traffic and bandwidth consumption used in maintaining consistent routing tables between routers.

Because BGP does not interact well with IGP protocols, Internal BGP (IBGP) Intra-AS Routing is implemented. When IBGP is used, BGP routes are not propagated into the IGP. Instead, all routers in the autonomous system use IBGP to communicate to each border router. Each router maintains two routing tables — one for internal routes and one for external routes.

Additionally, the routers support BGP-OSPF interaction, which permits importing BGP routes into OSPF. Support of BGP, IBGP Intra-AS Routing, and BGP-OSPF interaction ensures communications between a wide variety of dissimilar autonomous systems.

Additionally, BGP-4 provides the ability to configure the routing policies required by the Internet providers, (e.g., route aggregation, including aggregation of AS paths). This enhances route selection control. BGP-4 can coexist with RIP, EGP, OSPF, and static routing.

Bay Networks routers support BGP-3 as well as EGP-2. However, with the advent of BGP-4, BGP-3 has been moved to historical status and therefore should only be used if absolutely necessary.

Exterior Gateway Protocol Version 2 (EGP-2) EGP-2 is the exterior gateway protocol that features a neighbor acquisition mechanism that allows two routers to agree to support a mutual connection and exchange routing information. The EGP routing table contains a list of routers, the networks those routers can reach, and their associated cost metric. To maintain network reachability information using EGP, a router transmits its entire routing table in response to a poll command. Bay Networks routers support polling intervals from 120 to 480 seconds.

BOOTP Gateway (BOOTP Pass-Thru) Bay Networks routers support RFC 951 Section 8 of the Bootstrap Protocol (BOOTP) specification and RFC1542, Clarification's and Extension to BootP for DHCP.

With BOOTP Gateway, a Bay Networks router can transfer BOOTP packets, enabling diskless clients to boot from a server located on a network several hops away. BOOTP Gateway support can be enabled on individual network interfaces to receive and forward both BOOTREQUEST and BOOTREPLY packets to their destinations.

Bay Networks support for RFC1542 also ensures full routing support for IP networks using the Dynamic Host Configuration protocol for dynamic IP host address assignment.

Router Discovery Support of IP's Router Discovery function enables hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers. An extension of the Internet Control Message Protocol (ICMP), Router Discovery, eliminates the need for manual configuration of router addresses and is independent of any specific routing protocol.

Using Router Discovery, each router periodically multicasts discovery messages, referred to as Router Advertisements, from each of its multicast interfaces, announcing the IP address(es) of that interface. Hosts discover the addresses of their neighboring routers simply by listening for advertisements.

When a host attached to a multicast link initializes, it can multicast a Router Solicitation to ask for immediate advertisements rather than waiting for the next periodic ones to arrive. If no advertisements are forthcoming, the host may retransmit the solicitation. Any routers that subsequently initialize, or that were not discovered because of packet loss or temporary link partitioning, are eventually discovered by reception of their periodic (unsolicited) advertisements.

The router discovery messages enable hosts to discover the existence of neighboring routers, but not which router is best to reach a particular destination. If an inadequate first-hop router for a particular destination is chosen, the router issues an ICMP Redirect to the hosts. The Host Table will then consist of a default and ICMP-learned host routes.

Router advertisements include a user-configurable advertising rate and a lifetime field. The advertising rate specifies the frequency with which a router advertises its address. The lifetime field specifies the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts, in the absence of further advertisements. This is used to ensure that hosts eventually forget about routers that fail, become unreachable, or stop acting as routers.

Address Resolution Protocol (ARP) and Proxy ARP ARP enables an IP host to determine the MAC-layer address of a target host when all that is known is the target host's IP address.

To transmit a datagram, an IP device must know the IP destination address represented in the datagram's IP header. A router makes packet forwarding decisions based on the IP destination address. Once a routing decision has been determined, the router forwards the packet to the next hop network providing the best path to the packets ultimate destination. To accomplish this, the MAC-layer address of the next hop interface is added to the datagram and the packet is forwarded out the appropriate router interface.

If the next hop MAC-layer address is not known, the router must first broadcasts an ARP Request packet to determine the MAC-Address of the next hop interface. When the destination with the matching IP address receives the broadcast, it responds with its MAC-layer address, which is entered in the originating router's cache for future use.

Per RFC 826, a router and host must be attached to the same network segment to accomplish ARP. ARP Request broadcasts cannot be forwarded by another router to a different network segment. If a host requests the hardware address of a host on another network, Proxy ARP must be used. Proxy ARP support allows an intermediate router to answer the request for the remote destination host.

Reverse Address Resolution Protocol Server (RARP Server) An RARP server allows hosts to obtain IP addresses from the router. Hosts added to the network broadcast a RARP Request, specifying itself as the source and supplying its MAC-layer address in the frame's Destination Hardware Address field. When the RARP server receives the RARP request, it enters an IP address in the RARP request's Destination IP address field, changes the message type to a Reply, and sends the packet back to the host that transmitted the request, using the host's MAC-layer address.

Simple Network Management Protocol (SNMP) SNMP is the standard protocol used to monitor and control IP routers and attached networks. This transaction-oriented protocol specifies the transfer of structured management information between SNMP managers and agents. An SNMP manager, residing on a workstation,

issues queries to gather information about the status, configuration, and performance of the router. An SNMP agent, operating in each Bay Networks router, responds to the queries issued by the manager and generates activity reports. In addition to responding to SNMP queries, the router's SNMP agent software transmits unsolicited reports, referred to as traps, to the SNMP manager when events, such as the number of network collisions, exceed user-configured thresholds.

Each router maintains a set of configuration and performance variables in a database, referred to as a Management Information Base (MIB). All Bay Networks routers contain a MIB II-compliant SNMP agent that responds to SNMP SET/GET requests for configuration, monitoring, and control of network interfaces. Site Manager, Bay Networks node management application for Bay Networks routers, uses SNMP for router configuration, monitoring, and control. The routers may also be managed from popular general-purpose SNMP management systems such as Hewlett-Packard's OpenView, SunNet manager, and IBM's NetView for AIX.

Virtual Terminal Protocol (Telnet) Bay Networks enhances router installation and maintenance by supporting Telnet, the simple remote terminal protocol. Through incoming Telnet sessions, a Bay Networks router's Technician Interface (TI) can be accessed by a local or remote terminal. Outbound Telnet support enables TI to also originate an outgoing Telnet session to another Bay Networks router or to other network equipment that accepts inbound Telnet. This is used to access remote routers in nonroutine situations when Site Manager or SNMP is unavailable. Each instance of the TI supports a single outbound Telnet session.

The Technician Interface is based on a simple command line interpreter and provides SNMP-based access to the MIB, displays the router's event log, and supports file system management and other administrative commands.

Trivial File Transfer Protocol (TFTP) A Bay Networks router's support of TFTP allows a network management station to download configuration information to a router or group of routers and retrieve information from a router via Site Manager. Bay Networks routers include client and server implementations of TFTP, enabling efficient transmission and receipt of files across the internetwork. TFTP provides file transfer capabilities with minimal network overhead. Although TFTP uses UDP to transport files between network devices, it supports timeout and retransmission techniques to ensure data delivery.

File Transfer Protocol (FTP) The Bay Networks router's support of FTP enables a network management station to initiate router-to-host, host-to-router, and router-to-router data transfers over TCP via Site Manager. This implementation supports RFC 959 (File Transfer Protocol) to ensure that data is transferred reliably and efficiently. FTP is supported on all Bay Networks routers and by all the router's LAN, serial, and ATM interfaces.

Revised Internet Protocol Security Option (RIPSO) The IP implementation supports the Department of Defense (DOD) RIPSO on a per-interface basis. This support ensures that the integrity of datagrams requiring a high level of security is not compromised when received or transmitted by a Bay Networks router. RIPSO enables hosts to add security labels to IP datagrams for classification purposes. Through RIPSO, a host can label individual IP datagrams with one of four security classifications — Top Secret, Secret, Confidential, and Unclassified — and a set of protection authorities. These security labels can be compared on inbound, originated, or forwarded IP datagrams.

Bay Networks IP Enhancements

Bay Networks provides many advanced features as part of its IP implementation for its routers. These include IP Routing Policies, IP Multicast Routing support, Static Routes, Adjacent Hosts, Circuitless IP Interface Support, Configurable RIP Timers, Configurable Triggered Updates, Source Route Bridge End Station Support, NetBIOS over IP Support, Ping MIB, Crypto Resynchronization Support, and Blacker Front-End Support. These enhancements optimize internetwork reliability, availability, performance, and security.

IP Routing Policies IP Routing Policies are the rules that allow for the definition of criteria for routes accepted into domain and for routes advertised to other domains. The IP Policies govern the addition of routing information to the routing tables and propagation of routing information. The Policy Filter features "filterable" fields which are consistent among the different protocol and the syntax can represent single network entries and/or ranges of networks. Policy Filters also provide the ability to aggregate and deaggregate routes. Within Policy Filters, Import Route Filters are referred to as Accept Policies, and Export Route Filters as Announce Policies.

Announce Policy rules contain a parameter called the network advertisement list. This list controls the actual network advertisements that the router delivers to its neighbors. For BGP-4 and OSPF, it provides the ability to aggregate subnets and networks into supernets. Bay Networks recommends using the more powerful Policies.

Figure 5 | **Configurable RIP Interface Costs**

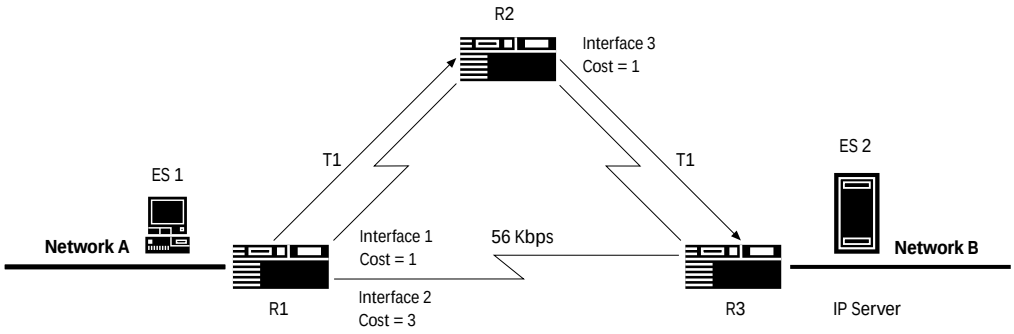

Configurable RIP Interface Costs The IP implementation supports configurable RIP costs on a per-interface basis. This feature is especially useful in topologies having two or more paths, of different bandwidths, connecting two networks. A path's cost is assigned during initial configuration, and can be changed anytime to meet new requirements.

For example, in Figure 5, traffic generated by end system ES 1 on Network A can be directed to the server on Network B over the two T1 circuits rather than the single 56-Kbps circuit with only "two hops." This can be accomplished by assigning a RIP interface cost of 3 to interface 2 of router R1, and assigning a RIP interface cost of 1 to interfaces 1 and 3 of routers R1 and R2, respectively. Because the total cost to reach router R3 via the T1 links is 2, compared to a cost of 3 for the 56-Kbps link, the primary path consists of T1 links. If the T1 network should fail, the lower bandwidth 56-Kbps link is automatically used as a backup. The ability to assign different cost values to each interface optimizes bandwidth use.

IP Multicast Routing Support By supporting IP Multicasting, a Bay Networks router allows messages to be sent to members of a multicast group simultaneously. Data packets are only sent to the end stations specified as group members. This improves bandwidth efficiency by reducing network traffic.

IP multicast routing is supported with the Distance Vector Multicast Routing Protocol (DVMRP), an early version of which is specified by the IETF RFC 1075. Based on distance vector or Bellman-Ford technology, DVMRP routes multicast a datagram within a single autonomous system. DVMRP also specifies the tunneling of IP multicasts through non multicast routing-capable IP domains. DVMRP can coexist with BGP, OSPF, EGP, RIP, and Static Routes. IP multicasting uses the Internet Group Management protocol (IGMP) as the protocol to communicate between hosts and multicast routers on a single physical network to establish a host's membership in particular multicast groups. IGMP also allows end stations to join and leave multicast groups.

Static Routes An administratively configured IP route can be manually entered into an IP routing table through the static routing feature. Available bandwidth is increased by eliminating the need to transmit dynamic routing updates, such as RIP and OSPF, over the network. Additionally, because static routes do not "age-out" of IP routing tables, remote offices or mobile workgroups using a dial-up service such as Dial-on Demand to communicate with a central site, are ensured that a data path exists.

Adjacent Host Support Adjacent Host Support allows a transmission path to be specified from a router to a host that resides on a locally attached network segment. This feature is typically configured for hosts that do not implement ARP and predefines the IP/data link address pair for each such local host. By supporting topologies that include non-ARP devices, the Adjacent Host feature enhances connectivity.

Adjacent hosts can also be configured for a local host that does support ARP to preempt the ARP process. By pre-resolving the host's IP/data link address pair, Adjacent Host configurations reduce network overhead by avoiding ARP handshaking.

Circuitless IP Interface The Circuitless IP Interface feature allows a backup IP address to be specified for a router without mapping it to a specific circuit. This ensures that the router is reachable if one or more of the router's IP interfaces becomes disabled. A malfunctioning router can still receive routing update messages and communicate with network management systems using its circuitless IP address, reducing the impact of hardware malfunction. IP traffic is received from and transmitted to the circuitless interface using the same method as any other IP interface.

Configurable RIP Timers Configurable RIP timers encompass three value added features — Configurable Broadcast Timers, Configurable Time Out Intervals, and Configurable RIP Hold Down Timers. These provide direct control over the amount of traffic generated on network links and are particularly beneficial in dial on demand environments as they reduce bandwidth consumption and eliminate the risk of dial connections being established due to propagation of RIP maintenance traffic.

- Configurable Broadcast Timers control the frequency in which a Bay Networks router propagates IP RIP routing table broadcasts. RIP broadcast intervals are now configurable from as low as once every five seconds to high as once every 24 hours.
- Configurable Time Out Intervals are used for entries within an IP RIP routing table. Standard IP RIP age out intervals will age out a route entries within the IP RIP routing table if an update for that route has not been received for 90 seconds (3 times the standard broadcast interval). Configurable age out intervals provides users direct control over how long RIP entries are maintained before being deleted if no updates are received.
- Configurable RIP Hold Down Timers enable the time period a router will wait before propagating network topology change information to be defined. This allows the router to group multiple changes into a single topology change update.

Configurable Triggered Updates The Configurable RIP Triggered Updates feature allows a Bay Networks router to be configured to either immediately propagate network change information when learned or suppress the information. Two configuration options are provided: Enabling triggered updates results in the router sending out network change information as soon as it is learned, while disabling triggered updates configures the router to send network change information only when a periodic RIP broadcasts is propagated.

Source Route Bridge End Station The Source Route Bridge End Station feature enables routable traffic generated in a source route bridge environment to be routed to end stations on remote LANs over a multiprotocol backbone. This reduces source route bridge overhead on a wide area network and maximizes network availability by rapidly rerouting around a failed link.

When Source Route Bridge End Station is enabled, a Bay Networks router attached to a Token Ring in a source route bridge environment functions as an end station and router. All traffic is source route bridged within the local Token Ring environment. IP traffic intended for a destination on a LAN interconnected via a multiprotocol backbone is routed over the backbone by the Bay Networks node.

NetBIOS over IP Support Bay Networks routers can route NetBIOS frames encapsulated within IP data-grams to provide efficient routing of NetBIOS information across an IP-based inter network. The router's NetBIOS over IP support is based on RFC 1001 and RFC 1002 broadcast node (b-node). This allows the router to rebroadcast NetBIOS packets beyond a local subnet to ensure unique NetBIOS name registration and provide immediate visibility of new NetBIOS resources as they become available on the network.

The router also provides a number of enhancements that improve the efficiency of routing NetBIOS over IP — NetBIOS Name Caching, NetBIOS Broadcast Filters, and NetBIOS Local Acknowledgment. These features improve network performance, enhance traffic control, and increase bandwidth availability.

Ping MIB The Bay Networks router's Ping MIB enables network availability and response time to be tracked. This provides a foundation for availability and response time reporting. Matrices such as those for source and destination IP addresses can be easily created using Site Manager, Bay Networks node management application for Bay Networks routers. Additionally, multiple destination addresses are supported for each source address.

Crypto Resynchronization A Bay Networks router can automatically detect the loss of synchronization between KG84A cryptographic encryption devices and initiate resynchronization. This capability enhances efficiency and availability by maintaining the communications link between two KG84A encryption devices. KG84A devices communicate over a point-to-point serial line and connect to a Bay Networks router via a V.35 Synchronous interface.

Blacker Front-End (BFE) Bay Networks routers can be connected directly to Blacker Front-End encryption devices to protect sensitive data transmitted over an unsecured X.25 network. The Blacker Front-End device provides the router with encryption services and access to the X.25 network. The Bay Networks router communicates with the Blacker Front-End Device over an X.25 Synchronous interface, which supports data rates between 1200 bps to 64 Kbps and complies with the 1983 DDN X. Host Interface Specification.

Local Area Network Support

All Ethernet, Token Ring, and FDDI network interfaces for Bay Networks routers support IP. The routers support SNAP and Ethernet encapsulation over Ethernet/802.3, SNAP encapsulation over FDDI, and LLC over Token Ring/802.5 media.

Wide Area Network Support

All serial interfaces for Bay Networks routers support IP. Serial interfaces operate at rates ranging from 1200 bps to 52 Mbps, full-duplex, and support V.35, RS232, RS449/RS422 balanced, X.21, MCT1/E1, ISDN BRI, and HSSI. The Synchronous interfaces support either internal or external clocking. Networks can also be interconnected via a variety of WAN services, including X.25, Frame Relay, SMDS, ATM DXI, ISDN, or point-to-point circuits using PPP or HDLC encapsulation.

Dial Back-up, Dial-on-Demand, and Bandwidth-on-Demand are also supported by the IP implementation over V.35 and RS232 interfaces.

ATM Network Support

Bay Networks router ATM link module interfaces and ATM Data Exchange Interface (DXI) software support Bay Networks IP implementation. ATM link modules operate at up to 155 Mbps and support SONET/SDH single and multimode fiber, and TAXI multimode fiber.

The ATM DXI operates over HSSI, V.35, and RS449 interfaces at up to 52 Mbps and can be used in all Bay Networks routers. This software interface fully complies with Modes 1a, 1b, and 2 of the ATM Forum's DXI specification for communications between routers and DSU/CSUs.

Traffic Management

Comprehensive traffic management capabilities are provided through Multiline Circuit Support, Uniform Traffic Filters, Traffic Prioritization, and Data Compression.

Multiline Circuits Multiline Circuits allows a single circuit to be composed of up to 16 individual serial network data paths, ensuring circuit availability in the event of a single data path failure. Multiline Circuits also increases bandwidth between two sites without the circuit management complexities associated with multiple circuits. Following initial configuration, the use of multiple data paths to form a single circuit is completely transparent.

Multiline Circuits provides two methods for transmitting traffic over its data paths — address-based selection and random selection. Address-based selection determines the path a packet takes based on its source and destination addresses. Once a path has been established for a given address pair, subsequent packets follow the same path. This ensures the sequentiality of packets and is a valuable feature for protocols that cannot tolerate packets received out of order.

Random selection determines the path each packet takes based on a randomly assigned number, which corresponds to a particular data path in the circuit. This provides for even distribution across the circuit to avoid congestion and is intended for use with protocols that can accept packets received out of sequence. The ability to select the method of transmitting data across the circuit enables Multiline Circuits to maximize the performance of a wide range of applications.

Uniform Traffic Filters Uniform Traffic Filters enables inbound and outbound traffic filters to be easily established for all network and bridge protocol traffic. Uniform Traffic Filters provides an efficient method for developing an effective and comprehensive network security strategy. In addition, Uniform Traffic Filters preserves WAN bandwidth and can increase performance by reducing network congestion.

Inbound traffic filters can be configured to accept or drop incoming packets from any Bay Networks router's local area, ATM, or serial network interface. Outbound traffic filters can be configured to drop outgoing packets destined for any Bay Networks router's serial interface. Additionally, Uniform Traffic Filters can be configured to execute a log action when a datagram's fields match the values defined in the filter.

Filters can be created using predefined protocol-specific fields or user-defined fields. Up to 31 inbound filters and 31 outbound filters (including Traffic Prioritization filters) can be defined for each protocol on every supported network interface. Filter precedence can be configured on an interface, reducing filter definition complexity. All filters are configured via Site Manager, the node management application for Bay Networks routers.

Traffic Prioritization Traffic Prioritization filters can assign a high priority to time-sensitive and/or mission-critical traffic, thereby reducing the occurrence of session timeouts and improving application response times. Priority filters can be configured that place packets into one of three priority queues — high, normal, or low — for transmission through a Bay Networks

router's outbound serial interface. Priority filters can be applied to the complete family of network and bridging protocols supported by Bay Networks routers.

Priorities can be assigned to packets based on their protocol, source network, destination network, packet type, and other protocol-specific fields, as well as other fields that are identifiable by an offset in a packet. The number of priority filters defined for a protocol on an interface depends on the number of outbound Uniform Traffic Filters assigned to the protocol on the interface. For example, if there are no outbound Uniform Traffic Filters defined for a protocol, then 31 priority filters can be assigned. However, if, for example, 16 outbound Uniform Traffic Filters are defined for a protocol, then only 15 priority filters can be assigned.

Traffic Prioritization can be configured to use either a strict dequeuing algorithm or a bandwidth allocation dequeuing algorithm to transmit packets across a serial line. Bay Networks strict dequeuing algorithm transmits all packets from the high-priority queue before transmitting packets from the normal and low-priority queues. The bandwidth allocation dequeuing algorithm allows packets from the normal and low-priority queues to be transmitted when the high-priority queue still contains packets, based on user-assigned bandwidth allocation percentages for each queue. This ensures that packets assigned lower priorities are transmitted in environments with large amounts of high-priority traffic. Each serial line attached to a Bay Networks router can use the strict or bandwidth allocation dequeuing algorithm and can be reconfigured anytime in response to changes in configuration and/or performance requirements.

Data Compression Configurable on a per-circuit or link basis, Bay Networks software-based Data Compression feature is supported by all Bay Networks routers, maximizing internetwork performance by reducing the amount of bandwidth required to transport LAN protocols over the wide area. Data Compression is currently supported over Dial-up lines, including ISDN, and leased lines using PPP. Based on a Lempel-Ziv algorithm, Bay Networks payload compression mechanism provides an aggregate compressed throughput of up to 512 Kbps, full-duplex, over links operating at up to fractional T1/E1 speeds.

Support is provided for either Continuous or Packet-by-Packet compression modes. Continuous mode maintains a compression history across packet boundaries, and requires that the histories at each end of the link be synchronized through the use of a reliable data link protocol. Packet mode resets the history for each packet, and does not require a reliable data link protocol. Continuous mode is recommended for maximum compression efficiency.

Network Management

Bay Networks offers comprehensive router and network management capabilities to ensure the efficient operation of mission-critical internetworks. Features increase diagnostic capabilities across the internetwork, simplify node and network configuration management, and interoperate with third-party solutions for increased interoperability.

Optivity/Internetwork Optivity/ Internetwork integrates Site Manager, Bay Networks node management application for Bay Networks routers with RouterMan™, an intuitive router monitoring application, and PathMan™, a graphical network diagnostic tool to simplify and improve management of complex router internetworks.

RouterMan offers real-time router performance and status reporting. The application's intuitive graphical user interface provides at-a-glance overall router status. Fault and performance statistics, history and analysis are provided for the overall router, by protocol or by interface. A simple color-coded interface and fault history log provides proactive detection and indication of potential router problems.

PathMan dynamically determines the complete data path between any two network end stations, assisting network managers in troubleshooting large complex networks. All network components on the selected path appear automatically in a system-generated display of the route, showing exactly how devices are connected within the network. Color-coded icons display each device's status, enabling rapid problem identification.

Site Manager is a platform-independent, SNMP-based application developed expressly for simplifying the configuration and management of Bay Networks routers. It provides an intuitive point-and-click user interface that streamlines the configuration process and eliminates cryptic commands. Site Manager offers central configuration management that simplifies network setup and expansion, real-time operations and monitoring, and real-time event and fault monitoring for efficient problem identification and isolation.

Site Manager is available for MS Windows, Sun SPARC, HP/9000 and IBM RS/6000 platforms.

Optivity/Internetwork operates independently or with the leading SNMP Platforms — HP OpenView, IBM NetView for AIX, and SunNet Manager for additional capabilities.

Standards

The IP implementation described in this information bulletin supports major IETF RFCs (see Table 1).

Table 1 | **IETF RP RFC Support**

RFC Number	Description
768	User Datagram Protocol (UDP)
783	Trivial File Transfer Protocol (TFTP)
791	Internet Protocol (IP)
792	Internet Control Message Protocol (ICMP)
793 and 1323	Transmission Control Protocol (TCP)
826	Address Resolution Protocol (ARP)
854	Virtual Terminal Protocol (Telnet)
877 and 1356	IP over X.25 Networks
903	Reverse Address Resolution Protocol (RARP)
904	Exterior Gateway Protocol (EGP) Version 2
950	Internet Subnetting Procedures
951	Bootstrap Protocol (BOOTP)
1001	Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Concepts and Methods
1002	Protocol Standard for a NetBIOS Service on a TCP/UDP Transport: Detailed Specifications
1009	Internet Gateway Requirements
1042	IP over IEEE 802 Networks
1058	Routing Information Protocol (RIP)
1063	Maximum Transmission Unit Discovery Option
1075	Distance Vector Multicast Routing Protocol (DVMRP)
1084	BOOTP Vendor Extensions
1108	Revised Internet Protocol Security Option (RIPSO)
1112	Internet Group Management Protocol
1155	Structure and Identification of Management Information
1156	Internet Management Information Base
1157	Simple Network Management Protocol (SNMP)
1188	IP over FDDI
1247	Open Shortest Path First (OSPF) Version 2
1256	Router Discovery
1267	Border Gateway Protocol (BGP) Version 3
1519	Classless Inter-Domain Routing (CIDR)
1542	Clarification's and Extension to BootP for DHCP
1654	BGP Version 4

MIB Information

The IP MIB defines a number of "objects" or variables to be monitored (see Table 2).

Table 2 | **IP MIB**

Object	Description
1. IP_create/delete	Indicates whether the IP interface has been created (1) or deleted (2)
2. IP_enable/disable	Indicates whether the IP interface has been enabled (1) or disabled (2)
3. IP_state	Indicates current state of IP: up (1), down (2), initialized (3), invalid (4),or not present (5)
4. IP_addr	Identifies the IP address this entry's addressing information pertains to
5. IP_interface_circuit	Identifies the circuit number that this interface operates over
6. IP_interface_mask	Identifies the subnet mask associated with the IP address of this entry
7. IP_interface_cost	Indicates the cost associated with this IP interface
8. IP_interface_cfg_bcastaddr	Identifies the specified IP broadcast address used for sending datagrams on this interface
9. IP_interface_bcastaddr	Identifies the broadcast address for sending datagrams on this interface
10. IP_ interface_mtu_discovery	Indicates whether the MTU discovery option is on (1) or off (2)
11. IP_interface_amr	Indicates whether the address mask reply is on (1) or off (2)
12. IP_interface_addr_res_type	Indicates which address resolution type is being used: ARP (1), probe (2), X.25 DDN (3), X.25 PDN (4), INARP (5), or ARPINARP (6)
13. IP_interface_asb	Indicates whether all subnet broadcasts are accepted and transmitted from this interface
14. IP_interface_proxy	Indicates whether the interface has proxy ARP on (1) or off (2)
15. IP_interface_host_cache	Indicates whether the host cache (address aging) is off (1) or states aging time; range is from 120 to 1200 seconds
16. IP_interface_udp_xsum	Indicates whether the UDP checksum verification is on (1) or off (2)
17. IP_interface_cfg_mac_addr	Identifies the user-configured MAC address of interface
18. IP_interface_mac_address	Identifies the actual MAC address of the interface
19. IP_interface_reasm_max_size	Indicates the size of largest IP datagram that can be reassembled
20. IP_interface_max_info	Indicates the maximum size of the non-MAC info field
21. IP_interface_in_receives	Indicates the total number of IP datagrams received from interface
22. IP_interface_in_hdr_errors	Indicates the number of input datagrams discarded due to error in header
23. IP_interface_in_addr_errors	Indicates the number of input datagrams discarded because of an invalid IP address
24. IP_interface_forw_datagrams	Indicates the number of input datagrams forwarded by router
25. IP_interface_in_unknown_proto	Indicates the number of locally addressed datagrams discarded because of an unknown or unsupported protocol
26. IP_interface_in_discards	Indicates the number of input datagrams discarded due to lack of buffer space
27. IP_interface_in_delivers	Indicates the total number of input datagrams successfully delivered
28. IP_interface_out_requests	Indicates the total number of IP datagrams supplied in response to requests for transmission
29. IP_interface_out_discards	Indicates the number of output IP datagrams discarded due to lack of buffer space

Table 2 | **IP MIB (continued)**

Object	Description
30. IP_interface_out_no_routes	Indicates the number of IP datagrams discarded because no routes could be found
31. IP_interface_reasm_timeout	Indicates the maximum number of seconds an interface can hold received fragments before it reassembles the message
32. IP_interface_reasm_reqds	Indicates the number of IP fragments received
33. IP_interface_reams_ok	Indicates the number of IP datagrams successfully reassembled
34. IP_interface_reasm_fails	Indicates the number of failures detected by the IP reassembly algorithm
35. IP_interface_frags_ok	Indicates the number of input datagrams successfully fragmented
36. IP_interface_frag_fails	Indicates the total number of input datagrams that could not be fragmented
37. IP_interface_frag_creates	Indicates the number of IP fragments generated
38. IP_interface_icmp_in_msgs	Indicates the total number of ICMP messages
39. IP_interface_icmp_in_errors	Indicates the number of ICMP messages with errors
40. IP_interface_icmp_in_dest_unrc	Indicates the number of ICMP "destination unreachable" messages received
41. IP_interface_icmp_in_param_prb	Indicates the number of ICMP "parameter problem" messages received
42. IP_interface_icmp_time_exced	Indicates the number of ICMP "time exceeded" messages received
43. IP_interface_icmp_src_quench	Indicates the number of ICMP "source quench" messages received
44. IP_interface_icmp_redirects	Indicates the number of ICMP "redirect" messages received
45. IP_interface_icmp_echo	Indicates the number of ICMP "echo" messages received
46. IP_interface_icmp_echo_reps	Indicates the number of ICMP "echo reply" messages received
47. IP_interface_icmp_in_timestmp	Indicates the number of ICMP "timestamp" messages received
48. IP_interface_icmp_tmstmp_reps	Indicates the number of ICMP "timestamp reply" messages received
49. IP_interface_icmp_in_addr_msk	Indicates the number of address mask requests received
50. IP_interface_icmp_add_msk_rep	Indicates the number of address mask reply messages received
51. IP_interface_icmp_out_message	Indicates the total number of ICMP messages which this interface attempted to send
52. IP_interface_icmp_out errors	Indicates the number of ICMP messages not sent due to problems
53. IP_interface_icmp_out_dest_unr	Indicates the total number of ICMP "destination unreachable" messages sent
54. IP_interface_icmp_out_time_exc	Indicates the number of ICMP "time exceeded" messages sent
55. IP_interface_icmp_out_parm_pb	Indicates the total number of ICMP parameter problem messages sent
56. IP_interface_icmp_out_src_qunc	Indicates the total number of ICMP "source quench" messages sent
57. IP_interface_cache_removes	Indicates the number of networks which have been flushed from cache due to aging
58. IP_interface_cache_networks	Indicates the total number of entries in the cache

Operation

A set of IP-specific parameters must be configured for each interface supporting IP (see Table 3).

Table 3 | **IP Configuration Parameters**

Parameter	Function	Action
Enable	Enables/disables IP routing on this interface	Default is Enable; set to Disable to deactivate IP routing on this interface.
Subnet Mask	Specifies the network and subnet portion of the 32-bit IP address	Enter the subnet mask for the class of the network connected to this interface in dotted decimal notation.
Broadcast Address	Specifies the broadcast address that the IP router uses to broadcast packets	Set to 0 to configure the IP router to use an all-1s address; optionally, enter the desired address in dotted decimal notation.
Interface Cost	Specifies the cost of the interface	Default is 1; optionally, enter value to 16.
MTU Discovery	Specifies whether the Maximum Transmission Unit reply option is enabled on the interface	Default is Off; set to On to enable this interface to respond to probe MTUs.
Addr Mask Reply	Specifies whether this interface generates ICMP address-mask-reply messages responding to valid request messages	Default is On; set to Off to disable ICMP address-mask-reply messages.
All Subnet Bcast	Specifies if the IP router floods received ASB datagrams across this interface	Default is On; set to Off to prohibit ASB flooding on this interface.
Address Resolution	Specifies whether this interface uses ARP to map 32-bit IP addresses to 48-bit Ethernet addresses	Default is Enable; set to Disable to disable ARP on this interface.
Proxy	Specifies whether this interface uses proxy ARP to respond to ARPs for a remote network	Default is Off; set to On to enable Proxy ARP on this interface.
Host Cache	Specifies whether the IP router ages entries in the interface's address-resolution cache and specifies the aging interval in seconds	Default is 1 (Off); optionally set to 120, 180, 200, 240, 300, 600, 900, or 1200.
Checksum	Specifies whether UDP checksum processing is enabled on this interface	Default is On; set to Off to disable UDP checksum processing.
MAC Address	Specifies a MAC address for this interface	Enter 0 to have router use its IP address and circuit's MAC address.

System Requirements

Bay Networks IP implementation described in this information bulletin is currently included in software Version 8.10 for the Bay Networks Access Node (AN™), Access Node Hub (ANH™), Access Feeder Node (AFN™), Access Stack Node (ASN™), Feeder Node (FN™), Link Node (LN™), Concentrator Node (CN™), Backbone Link Node (BLN™), and Backbone Concentrator Node (BCN™).

Ordering Information

IP is available in a variety of software suites for the Bay Networks AN, ANH, ASN, BLN, BCN, LN, and CN (see Table 4).

Table 4 | **Ordering Information**

Model Number	Description
42032V###*	IP Access software suite for AN/ANH (includes IP) — 2 MB Flash
42034V###*	IP Access software suite for AN/ANH (includes IP) — 4 MB Flash
40044V###*	System software suite for ASN (includes IP)
40032V###*	System software suite for BLN and BCN (includes IP)
40012V###*	System software suite for LN and CN (includes IP)
40022V###*	Corporate software suite for AFN (includes IP)
40062V###*	Corporate software suite for FN (includes IP)

* ### = Software version number (e.g., Version 8.10 = 081)

IP is also available in the Corporate software suite for the Bay Networks AN, ANH, ASN, BLN, BCN, LN, and CN. It is also available in the Remote Office software suite for the Bay Networks AN and ANH.

Bay Networks

For more sales and product information, please call **1-800-8-BAYNET**.

United States

Bay Networks, Inc.
4401 Great America Parkway
Santa Clara, CA 95054
Phone: 1-800-8-BAYNET

Bay Networks, Inc.
8 Federal Street
Billerica, MA 01821-5501
Phone: 1-800-8-BAYNET

Europe, Middle East, and Africa

Bay Networks EMEA, S.A.
Les Cyclades – Immeuble Naxos
25 Allée Pierre Ziller
06560 Valbonne, France
Fax: +33-92-966-996
Phone: +33-92-966-966

Intercontinental

Bay Networks, Inc.
8 Federal Street
Billerica, MA 01821-5501
Fax: 508-670-9323
Phone: 1-800-8-BAYNET

World Wide Web: **http://www.baynetworks.com**

6/95 DA591-1975EC-C

Printed on recycled paper