



Software Product Description

PRODUCT NAME: Distributed Routing Software, Version 1.1

SPD 54.96.03

This SPD describes Version 1.1 the Distributed Routing software, which provides multiprotocol routing and bridging. This software is implemented on the products listed in Table 1. Information applies to all products with the interface being discussed unless otherwise specified. In this SPD, the term "Router" refers to an implementation of the Distributed Routing software.

Table 1

Distributed Routing Software Version 1.1 Products

Product	Interfaces	Form Factor
RouteAbout Access EW	Ethernet, 2 T1/E1 Serial Lines	Stand-Alone, MultiStack or DEChub 90 Module
RouteAbout Access EI	Ethernet, Basic Rate ISDN, T1/E1 Serial Line	Stand-Alone, Multi- Stack, or DEChub 90 Module
RouteAbout Central EW	2 Ethernet, 8 T1/E1 Serial Lines	Stand-Alone in DEChub ONE or Full-Height DEChub 900 Module

Version 1.0A of the Distributed Routing software is described in SPD 54.96.02. Version 1.0A is implemented on the following products:

- RouteAbout Access TW
- DECswitch 900EE
- DECswitch 900EF
- DECswitch 900ET

DESCRIPTION

Overview

The Distributed Routing Software is available in two packages. The IP package includes IP routing along with bridging and all WAN services. The Multiprotocol package supports all of the above plus routing for these protocols:

- IPX
- AppleTalk
- DECnet Phase IV
- DECnet/OSI
- DLSw
- SDLC Relay

The Bridging function interconnects networks for protocols that are not routed by the Router, and for protocols that do not have a Network layer to support routing. The Router can route some protocols, while concurrently bridging others. All models support Transparent bridging as defined in the IEEE 802.1d protocol.

All of the RouteAbout models described in this SPD have at least one T1/E1 serial interface. The serial interface supports the PPP, Serial Line Protocol, and SDLC data links. The X.25 and Frame Relay public network services are supported. V.25bis is supported for dial backup. The Bandwidth Reservation System allocates the WAN bandwidth to classes of traffic, prioritizes traffic within each class, and also provides MAC filtering.

The RouteAbout Access EI has a single Basic Rate ISDN interface. This interface can serve as a backup for WAN restoral when the T1/E1 serial line fails, supporting full multiprotocol bridging and routing. The Basic Rate ISDN interfaces also support dial on demand to a separate site for each of the two B channels. For dial on demand, TCP/IP and OSI static routing is supported, as

well as IPX routing. IPX spoofing reduces the amount of acknowledgment traffic that traverses the WAN.

The Distributed Routing software is factory installed in the flash memory of the Brouter. Updates and reloads can be performed from a load host using the BOOTP/TFTP protocol. Configuration and management are through the Brouter command line interface, either locally or remotely via TELNET. SNMP Gets and Traps are supported for monitoring. Limited HUBwatch support is also provided.

TCP/IP Routing

Routing Protocols: The IP implementation routes data in accordance with the TCP/IP standards. Routing table entries may be static, in which case they are configured by the user from the console, or they can be dynamically created by routing protocols. The Open Shortest Path First (OSPF), Routing Information Protocol (RIP), Border Gateway Protocol (BGP-4), Exterior Gateway Protocol (EGP), and Integrated Intermediate System to Intermediate System (IS-IS) protocols are supported.

Subnetting: Subnetting support is fully compliant with RFC 950. Any number of IP networks may be subnetted. When using OSPF or IS-IS, the Brouter supports variable length subnet masks as specified in RFC 1009. Subnet masks are specified on a per-subnet basis. When using RIP, subnet masks are specified on a per-network basis. In this case, a given IP network may only have one subnet mask.

Fragmentation: If the destination network does not support packets as large as those to be sent, the Brouter fragments the packets before transmission.

Access Control: The IP implementation supports selective packet filtering for security. Access control lists can be applied separately to each interface for either incoming or outgoing traffic. Packets can be filtered based on the source or destination address, IP protocol number, or TCP/UDP port number.

BGP-4: The Brouter supports the Border Gateway Protocol (BGP-4) as defined in RFC 1745 and 1771. The BGP-4 implementation is not downwardly compatible with the previous versions of BGP. BGP-4 is an exterior gateway routing protocol used to exchange network reachability information among autonomous systems. An autonomous system is analogous to a domain—a collection of routers and end nodes administered by an organization.

The Brouter operates as a BGP speaker, exchanging routing information with other BGP speakers within the same domain and other domains. It propagates routing information learned via BGP into the routing tables of other IGP protocols.

OSPF: The OSPF implementation allows multiple "best routes" to a destination network. Load balancing is supported. OSPF loopback detects loopback and makes the lines unavailable for traffic transmission. Access controls are supported. The OSPF MIB is supported for SNMP monitoring.

MOSPF Multicast Routing: The IP implementation supports the routing of IP multicast datagrams, which are identified as packets whose destinations are class D IP addresses (where the first byte lies in the range 224-239 inclusive). A single multicast datagram may be delivered to multiple destinations, called a multicast group. Both Peer-to-Peer and Client/Server groups are supported. Multicast routing is achieved through Multicast Extensions to OSPF (MOSPF) using the point-to-multipoint interface type.

The Distributed Routing software includes several native multicast applications. The ICMP ping command in the IP console can use an IP multicast address. The software can also be configured to send SNMP traps to one or more IP multicast addresses.

DVMRP Routing: The Distributed Routing software includes the *mrouterd* routing daemon, allowing the Brouter to be used in the MBONE in addition to or instead of UNIX workstations. Also, support of the Distance Vector Multicast Routing Protocol (DVMRP) allows MOSPF domains to be substituted for collections of DVMRP tunnels, easing bandwidth demands.

The DVMRP/MOSPF implementation can be run in one of three modes, described below in order of increasing functionality. In Mode 1, the Brouter serves as a regular DVMRP router, like a UNIX workstation running the *mrouterd* program.

In Mode 2, the Brouter connects an MOSPF domain to the MBONE via one or more encapsulated DVMRP tunnels. In this mode, the Brouter advertises selected MOSPF networks on the MBONE DVMRP network and advertises a subset of the DVMRP sources on the MOSPF networks as OSPF AS external LSAs. An MOSPF domain joined to the MBONE in this way receives the benefit of MOSPF's pruning, so that only those multicast datagrams with active group members are forwarded on the MOSPF domain.

In Mode 3, the Brouter serves as an MOSPF router, using an MOSPF domain as a "transit" network. In this mode, DVMRP runs over MOSPF, as if the entire MOSPF domain were a single LAN. This allows an MOSPF domain to replace a collection of DVMRP tunnels, decreasing multicast traffic.

ISDN Support: Dial on Demand with an ISDN B channel as the primary circuit is supported for IP static routing. Each B channel can be routed to a separate des-

mination. The two B channels cannot be aggregated to form a higher speed link to the same destination.

The Brouter can automatically establish an ISDN connection to the destination router when its serial line fails. In this WAN restoral mode, all multiprotocol bridging and routing is supported, just as on the serial line.

X.25 Support: IP traffic is supported on X.25 networks using either the encapsulation and addressing procedures specified in RFC 877 and RFC 1356, or those specified in DDN Standard X.25.

Frame Relay Support: The Brouter supports routing of IP traffic over Frame Relay networks as specified in RFC 1490. The Brouter Frame Relay interface also supports IP traffic over PPP via an encapsulation technique. ARP correctly resolves MAC addresses, while Inverse ARP maps the MAC address to the IP address, as defined in RFC 1293.

ICMP: The ICMP implementation generates ICMP error and informational messages in the appropriate circumstances.

IP Standards: The IP implementation is based on the following set of RFCs:

- RFC 768—User Datagram Protocol
- RFC 791—Internet Protocol
- RFC 792—Internet Control Message Protocol
- RFC 793—Transmission Control Protocol
- RFC 854—TELNET Protocol Specifications
- RFC 877—Transmission of IP Datagrams Over Public Data Networks
- RFC 888—"STUB" Exterior Gateway Protocol
- RFC 904—Exterior Gateway Protocol Formal Specifications
- RFC 919—Broadcasting Internet Datagrams
- RFC 922—Broadcasting Internet Datagrams in the Presence of Subnets
- RFC 925—Multi-LAN Address Resolution
- RFC 950—Internet Standard Subnetting Procedure
- RFC 951—Bootstrap Protocol
- RFC 1009—Requirements for Internet Gateways
- RFC 1058—Routing Information Protocol
- RFC 1112—Host Extensions for IP Multicasting
- RFC 1123—Requirements for Internet Hosts
- RFC 1191—Path MTU Discovery
- RFC 1195—Use of OSI IS-IS for Routing in TCP/IP and Dual Environments

- RFC 1247—OSPF Version 2
- RFC 1293—Inverse Address Resolution Protocol
- RFC 1340—Assigned Numbers
- RFC 1350—TFTP Protocol Version 2
- RFC 1356—Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode (X.25 Support Only)
- RFC 1360—IAB Official Protocol Standards
- RFC 1395—Bootstrap Protocol (BOOTP)
- RFC 1490—Multiprotocol Interconnect Over Frame Relay Networks
- RFC 1745—BGP-4/IDRP for IP—OSPF Interaction
- RFC 1771—A Border Gateway Protocol 4 (BGP-4)

Novell NetWare Routing

The IPX implementation routes NetWare traffic in accordance with the Novell specifications for the IPX protocols. A Multiprotocol software license is required for Novell NetWare routing.

IPX RIP and SAP: IPX support includes full implementations of NetWare Routing Information Protocol (RIP) and Service Advertising Protocol (SAP). The Brouter keeps multiple equal cost routes to a given remote IPX network but selects a primary route that it uses exclusively when the route is available. Path splitting is not supported. To conserve WAN bandwidth, the frequency of RIP and SAP periodic updates can be set to nonstandard values. IPX Type 20 Propagation is supported.

Filtering: The Distributed Routing software supports access controls, SAP filtering, and GNS response suppression.

LAN Encapsulation: The Distributed Routing software supports all of the Novell sanctioned encapsulations for IPX (see below). The Brouter supports only one type of encapsulation at a time on each interface.

- ETHERNET_II
- ETHERNET_8022 (default as of NetWare V4.0)
- ETHERNET_8023 (IEEE 802.3 without 802.2)
- ETHERNET_SNAP (used for bridging)
- TOKEN-RING MSB/LSB (DSAP/SSAP)
- FDDI (8022 DSAP/SSAP)
- FDDI_SNAP

Serial Line Support: The IPXWAN Control Protocol supports use of PPP for native IPX traffic according to RFC 1634. The options are not supported. IPXCP is also implemented according to RFC 1552, but the options are not supported. IPX also runs over the Serial Line Protocol.

IP Encapsulation Over WAN Links: Novell IPX is supported over wide area links via IP encapsulation according to RFC 1234.

ISDN Support: Dial on demand is supported for Novell NetWare routing. Each B channel can be routed to a separate destination. The two B channels cannot be aggregated to form a higher speed link to the same destination. Novell server spoofing is supported, where the router at the Novell server site responds to server KeepAlive messages on behalf of the clients. This spoofing eliminates a great deal of WAN traffic.

The Brouter can automatically establish an ISDN connection to the destination router when its serial line fails. In this failover mode, all multiprotocol bridging and routing is supported, just as on the serial line.

X.25 Support: Native IPX traffic is supported over X.25 in accordance with RFC 1356. The implementation uses the Call User Data (CUD) field for IPX.

Frame Relay Support: The Brouter supports routing of IPX traffic over Frame Relay networks as specified in RFC 1490. ARP correctly resolves MAC addresses, while Inverse ARP maps the MAC address to the IPX address, as defined in RFC 1293.

IPX Standards: The IPX implementation is based on the following specifications:

- IPX Router Specification
- Advanced NetWare V2.0 Internetwork Packet Exchange Protocol
- Advanced NetWare V2.0 Service Advertising Protocol
- NetWare Driver Specifications for Network Interface Cards
- Novell IPX Management Information Base
- Novell RIP-SAP Management Information Base
- RFC 1234—Tunneling IPX Traffic Through IP Networks
- RFC 1293—Inverse Address Resolution Protocol
- RFC 1356—Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode (X.25 Support Only)
- RFC 1490—Multiprotocol Interconnect Over Frame Relay Networks
- RFC 1552—PPP Internetwork Packet Exchange Protocol (IPXCP)
- RFC 1634—Novell IPX Over Various WAN Media (IPXWAN)

AppleTalk Routing

The Distributed Routing software routes AppleTalk traffic in accordance with the AppleTalk Phase 1 and Phase 2 protocols. These two protocols are implemented as two separate modules with their own configuration and console monitors. The phases have different encapsulations for each network which allows them to be run in parallel without conflict. A transition gateway provides connectivity between the two protocols for migration purposes.

Routing Table Maintenance Protocol: The Brouter maintains its AppleTalk routing tables using of the Routing Table Maintenance Protocol (RTMP). The Phase 2 implementation of RTMP has two extensions. The propagation of bad entries is speeded up with a notify neighbor technique. A second technique, split horizon, is used to shrink the size of the RTMP route update. Split horizon can be disabled to support AppleTalk Phase 2 in partially meshed Frame Relay networks.

Zone Information Protocol: The Brouter maintains a Zone Information Table (ZIT) through use of the Zone Information Protocol (ZIP). The ZIT consists of zone information associated with each network in the routing table. Phase 2 extends ZIP to allow zone lists for each network range along with a default zone name.

Name Binding Protocol: The Brouter participates in and facilitates the name binding process through the use of the Name Binding Protocol.

Network and Zone Filters: The AppleTalk Phase 2 implementation supports Network and Zonename filters for each interface. There are separate filter lists for incoming or outgoing information. The Brouter does not advertise filtered Zone information in the specified direction. Both inclusive (allowed Zone) and exclusive (blocked Zone) lists are supported.

Transition Gateway: The transition gateway provides connectivity between the AppleTalk Phase 1 and Phase 2, allowing Phase 1 AppleTalk nodes to connect to Phase 2 nodes.

EtherTalk Protocol: For AppleTalk Phase 1, the Brouter uses Apple's EtherTalk packet encapsulation for DDP packets transmitted on the Ethernet LAN. For Phase 2, it uses EtherTalk IEEE 802.3 encapsulation.

TokenTalk Protocol: For AppleTalk Phase 2, the Brouter uses Apple's TokenTalk encapsulation for DDP packets transmitted on 4 or 16 Mbps IEEE 802.5 Token Rings. Phase 1 AppleTalk is not supported on Token Ring.

Apple Address Resolution Protocol: The Apple Address Resolution Protocol (AARP), in conjunction with EtherTalk and TokenTalk, maintains hardware to protocol address mappings.

Serial Line Support: The Brouter encapsulates Phase 1 and Phase 2 AppleTalk on Serial Line Protocol serial lines using a proprietary method. The PPP protocol is supported only for Phase 2. The AppleTalk Control Protocol used for PPP adheres to RFC 1378. The only option negotiated is the AppleTalk address; all other options are rejected.

IP Encapsulation Over WAN Links: AppleTalk is supported over wide area links via IP encapsulation.

Half Router: The Brouter can act as an AppleTalk Half Router on PPP and Serial Line Protocol serial lines to support interoperability with other vendors' routers. By default, the Brouter is configured as a Half Router, meaning that it is not assigned an AppleTalk address. Both the network number and the AppleTalk node number are set to 0.

Frame Relay Support: The Brouter supports routing of AppleTalk Phase 1 and Phase 2 over Frame Relay networks as specified in RFC 1490. AppleTalk ARP correctly resolves MAC addresses, while Inverse ARP maps the MAC address to the AppleTalk address, as defined in RFC 1293. Manual configuration of AppleTalk Node and Network numbers onto specific DLCIs is also supported.

The Disable Split Horizon feature provides support for AppleTalk Phase 2 in partially meshed Frame Relay networks by ensuring that all routing tables are propagated from the "hub routers." A hub router is a router connected to two or more routers over a partially meshed network.

AppleTalk is not supported over X.25 except when encapsulated in IP.

AppleTalk Standards: The AppleTalk routing implementation conforms with the following specifications for Phase 1 and Phase 2:

- *Inside AppleTalk*, March 1989, Apple Computer, Inc., Addison Wesley (Phase 1)
- *Inside AppleTalk, Second Edition*, May 1990, Apple Computer, Inc., Addison Wesley (Phase 2)
- RFC 1243—AppleTalk MIB
- RFC 1293—Inverse Address Resolution Protocol
- RFC 1378—PPP AppleTalk Control Protocol (ATCP)
- RFC 1490—Multiprotocol Interconnect Over Frame Relay Networks

DECnet Phase IV Routing

The Distributed Routing software routes DECnet Network Architecture (DNA) packets in accordance with the Phase IV router specifications for a Level II router. Certain extensions from the Phase IV+ architecture are also supported.

The basic routing functions are provided, with network management through the Brouter command line interface from the local and remote Telnet consoles. Remote network management using the NICE protocol is not supported, nor is network booting via MOP. Access control facilities protect groups of nodes from other nodes in the network. An area routing filter controls which areas routing information will be accepted from, and which areas routing information will be sent for. Load balancing is not supported. A Multiprotocol software license is required for DECnet Phase IV routing.

Support for Use of Arbitrary MAC Addresses: The Distributed Routing software supports extensions to Phase IV that allow the use of arbitrary MAC Individual Station Addresses on Token-Ring networks. The Brouter supports both networks where arbitrary MAC addressing is used exclusively and networks where it is used in combination with standard Phase IV addressing.

Event Logging: The Brouter provides a large number of Phase IV diagnostic error messages. All of the appropriate Class 4 events (Routing) are logged through the error logging facilities. Additional messages report errors that are not specified in the DNA network management protocol and also provide tracing facilities.

Network Control Program: The DNA implementation can be fully configured using commands with a syntax modeled after Network Control Program (NCP) commands. Both SET and DEFINE operations are supported, as well as the corresponding SHOW and LIST operations.

X.25 Support: The Brouter routes DECnet Phase IV traffic using statically assigned X.25 routing circuits. DECnet Phase IV connections over X.25 are supported only to other Brouters.

Frame Relay Support: The Brouter supports routing of DECnet Phase IV traffic over Frame Relay networks as specified in RFC 1490.

DECnet Phase IV Standards: The DECnet Phase IV implementation is based on the following specification:

- *DECnet Digital Network Architecture Phase IV Routing Layer Functional Specification, Version 2.0.0*, December 1993
- RFC 1490—Multiprotocol Interconnect Over Frame Relay Networks

DECnet/OSI (Phase V) Routing

The DECnet/OSI (Phase V) implementation routes packets in accordance with the Phase V and Phase IV router specifications of the DECnet DNA protocol family. These in turn are based on the OSI protocols for Connectionless Mode Network Service (CLNP), Intermediate System to Intermediate System (IS-IS) routing, and End System to Intermediate System (ES-IS) routing. A Multiprotocol software license is required for DECnet/OSI routing.

Phase IV Compatibility: The implementation includes all of the Phase IV compatibility features that are part of the Phase V specifications, which allow interconnections of Phase IV and V hosts and networks. The Router can be configured to use the Phase IV or Phase V routing protocols at Level 1 (intra-area) and Level 2 (inter-area). All combinations of routing protocols at Level 1 and Level 2 are supported.

IS-IS and ES-IS Routing: The Router routes data in accordance with the Intermediate System to Intermediate System (IS-IS) and End System to Intermediate System (ES-IS) protocols. Routing table entries are created dynamically by the operation of these protocols. Load balancing is not supported.

ISO Function Support: The following ISO 8473 functions are implemented:

- PDU Composition
- PDU Decomposition
- Header Format Analysis
- PDU Lifetime Control
- Route PDU
- Forward PDU
- Segmentation
- Discard PDU
- Error Reporting
- PDU Header Error Detection

The following ISO 9542 functions are implemented:

- Report Configuration by Intermediate Systems
- Record Configuration by Intermediate Systems
- Configuration Notification
- Request Redirect
- PDU Header Error Detection

All mandatory functions of ISO 10589, with the exception of Timer Jitter, are implemented. In addition, the following optional functions are supported:

- PDU Authentication

- Subnetwork Dependent functions for Broadcast Subnetworks
- Subnetwork Dependent functions for ISO 8208 Subnetworks (not completely based on the standard)
- Subnetwork Dependent functions for point-to-point Subnetworks
- Level 2 IS-IS functions
- Reachable Address Prefixes

Integrated IS-IS Support: The Distributed Routing software supports Integrated Intermediate System to Intermediate System (IS-IS) routing in accordance with RFC 1195. RFC 1195 defines IP extensions to ISO 10589.

ISDN Support: Dial on Demand with an ISDN B channel as the primary circuit is supported for DECnet/OSI static routing. Each B channel can be routed to a separate destination. The two B channels cannot be aggregated to form a higher speed link to the same destination.

The Router can automatically establish an ISDN connection to the destination router when its serial line fails. In this failover mode, all multiprotocol bridging and routing is supported, just as on the serial line.

Frame Relay Support: The Router supports routing of DECnet/OSI traffic over Frame Relay networks as specified in RFC 1490.

X.25 Support: DECnet/OSI traffic is not supported on X.25 networks.

Network Management Support: Network Management is available from the local and Telnet remote consoles. Configuration is via the Distributed Routing Software command line interface. For tracing, the Trace Route diagnostic tool displays the hop-by-hop path between nodes. The RFC 1139 Echo function tests the responsiveness of target nodes and also generates echo replies, much like the TCP/IP ping function. Remote network management operations using NICE or CMIP are not supported, nor is network booting via MOP.

DECnet/OSI (Phase V) Standards: The DECnet/OSI implementation is based on the following specifications:

- *DECnet Digital Network Architecture (Phase V) Network Routing Layer Functional Specification*, Part Number EK-DNA03-FS-001, Version 3.0.0, July 1991.
- ISO 8473—Protocol for Providing the Connectionless-Mode Network Service
- ISO 9542—End System to Intermediate System Routing Exchange Protocol

- ISO 10589—Intermediate System to Intermediate System Intra-Domain Routing Information Exchange Protocol
- FIPS PUB 146—GOSIP, U.S. Department of Commerce, August 24, 1988
- RFC 1139—An Echo Function for ISO 8473
- RFC 1195—Use of OSI IS-IS for Routing in TCP/IP and Dual Environments
- RFC 1490—Multiprotocol Interconnect Over Frame Relay Networks

Data Link Switching (DLSw)

DLSw is supported on the Ethernet and serial ports of the products described in this SPD. The Distributed Routing software routes SNA wide area and local area data over TCP/IP networks as specified in RFC 1434. The SNA wide area traffic supported is SDLC, while the local area traffic supported is IEEE 802.2 Logical Link Control Type 2 (LLC2) over Token Ring or Ethernet. The Bandwidth Reservation system, described in detail below, provides the allocation of WAN bandwidth and prioritization needed to support DLSw traffic. The DLSw implementation interoperates with the IBM 6611, IBM 2210, Proteon CNX/DNX/RBX, and compatible products. The IBM 6611 Version 1 Release 3 implementation is the latest tested version. A Multiprotocol software license is required for DLSw.

The DLSw implementation supports T2.0 (PU Type 2.0) and T2.1 nodes, as well as Type 4 and 5 SNA Subarea nodes. T2.0 nodes (such as IBM 3174s) and T2.1 nodes (such as IBM AS/400 systems) can be connected to the Brouter via SDLC or LLC2. They can then communicate across the TCP/IP network to a host connected to another Brouter via LLC2. T2.1 nodes can communicate directly without host intervention via DLSw across a TCP/IP backbone.

The Brouter at the host site can be connected to a T4 front-end processor or directly connected to a T5 host equipped with a LAN to mainframe attach device such as an IBM 3172 Token Ring Controller. In either case, the LLC2 data link protocol is used for the connection. T4 to T4 (FID4) connections are not supported.

Instead of bridging SNA traffic, DLSw terminates the Level 2 SDLC or LLC2 data link, using TCP/IP as a reliable transport when both the source and destination are not on the same Brouter. Data link termination can decrease timeouts and retransmissions, while decreasing WAN traffic by eliminating SDLC polling on the WAN.

When using DLSw, SDLC attached devices appear to be attached via Token Ring to a 3745, 3172, or similar

device. The Brouter supports a large number of downstream SDLC stations by mapping a LAN address to each SDLC station.

All of the RFC 1434 functions of the DLSw protocol for SNA are supported. These functions include finding destinations with the CANUREACH protocol, coordinating the termination of SDLC and LLC2 sessions, reestablishing the logical link between two terminated link sessions, acting as a transport gateway using TCP/IP for data transfer, and providing for orderly disconnect. DLSw dynamically establishes TCP/IP sessions on demand as needed. DLSw can also automatically re-establish TCP connections upon power-up or after session loss. The size of the TCP receive buffer can be configured, improving performance over links with high delay, such as Frame Relay networks.

LLC2 Termination: The Brouter can support hundreds of concurrent locally terminated LLC2 sessions, as described above. DLSw uses a virtual segment (added to the RIF) on Source Routed networks.

SDLC Termination: The Distributed Routing software provides SDLC Primary Link Station functionality for DLSw. This allows the Brouter to communicate with SDLC Secondary Link Station devices. The Brouter supports SDLC termination of T2.0 (such as 3270) and T2.1 (such as APPC) devices, including SDLC multidrop support. T2.0 and T2.1 nodes can be mixed on the same multidrop link. Since the SDLC data link is terminated in the Brouter, no SDLC control information is passed between Brouters. The maximum number of SDLC stations supported is a function of memory available on the Brouter and configuration.

Local Conversion: The Distributed Routing software supports both SDLC and LLC2 termination concurrently in the same Brouter. It also supports the connection of locally attached SDLC devices to locally attached LLC devices.

Dynamic Discovery Via DLSw Groups: In addition to the static entry of DLSw router neighbors, DLSw neighbors can dynamically find each other by using the DLSw Group Membership functionality based on Multicast OSPF. DLSw groups alleviate the need for long lists of static IP addresses. The Brouter can be a member of as many as 64 groups. Both Client/Server and Peer-to-Peer groups are supported.

Flow Control: DLSw supports both TCP and DLSw session flow control. Two Distributed Routing software enhancements to DLSw, ENTER_BUSY and EXIT_BUSY, support DLSw circuit level flow control, allowing one end of the circuit to notify the other end of congestion and the clearing of congestion. In addition, DLSw will flow control the SNA data link with Receiver Not Ready (RNR) in the event of TCP transmit congestion. Brouters congested on TCP receive decrease their TCP

window. TCP flow control is accomplished using the TCP slow start algorithm.

Broadcast Reduction: The Brouter limits unnecessary broadcasts by using several caches. First, the Brouter caches the IP address of the destination DLSw router for all MAC addresses learned from ICANREACH messages. The size of the MAC address to IP address cache is configurable. Secondly, the Brouter caches all pending CANUREACH messages to avoid generating redundant messages. On the Token-Ring interface, the Brouter caches Route Information Fields (RIFs) to reduce the number of Explorer broadcasts.

Ethernet Support: DLSw supports LLC2 termination on Ethernet. Any Ethernet attached SNA devices such as 3174s, 3745s, and so on, are supported. In most topologies, only one DLSw Brouter can be attached to an Ethernet segment. Also, the topology must not include parallel bridge and DLSw paths between source and destination Ethernet end stations.

DLSw MIB: The Brouter supports Gets and Traps of DLSw MIB variables as defined in the IBM 6611 implementation.

DLSw Standards: The DLSw implementation is based on the following specifications:

- RFC 1434—Data Link Switching: Switch-to-Switch Protocol
- 6611 Network Processor Network Management Reference IBM Document GC30-3567

Boundary Access Node

The Boundary Access Node (BAN) function connects remote SNA sites to the central host site directly over Frame Relay. BAN connects multiple T2.0 and T2.1 end nodes over Frame Relay to the IBM Network Control Program (NCP Version 7 Release 3 or later) running on the front-end processor or host. Thus, dedicated SDLC links are replaced by Frame Relay connections. BAN eliminates the need for a single Frame Relay PVC per T2.0 or T2.1 end node. Also, there is no need for a router at the host site. The Boundary Access Node function works in two modes - bridge mode and DLSw mode, as described below. A Multiprotocol software license is required for the Boundary Access Node function.

Bridge Mode: BAN Type 1 is Bridge Mode. The Brouter acts as a bridge between the host/front-end processor and the end node. The Brouter converts received frames to IEEE 802.5 bridged frame format and transmits them across the Frame Relay network in accordance with RFC 1490.

DLSw Mode: BAN Type II is DLSw BAN. The Brouter terminates the LLC2 traffic received in accordance with the DLSw standard. At the same time, the Brouter establishes a virtual LLC2 connection over the Frame Relay network between the host/front-end processor and the end station. End stations on Token Ring, Ethernet, and SDLC links are supported.

SDLC Relay

The SDLC Relay function allows pairs of point-to-point Synchronous Data Link Control (SDLC) devices to be connected locally or remotely via an internetwork. This is done by connecting the SDLC device to a synchronous interface on the Brouter and then passing SDLC frames between the Brouters encapsulated in UDP/IP packets. The destination IP address is based on the source serial port of the frame. High-Level Data Link Control (HDLC) frames are supported in the same way. A Multiprotocol software license is required for SDLC Relay.

The operation of the Relay is transparent to the SDLC devices; the Relay looks like a full-duplex leased line. There is no processing of the data in the SDLC frames, other than protection from undetected data errors. The Relay does not support half-duplex or multi-drop operation.

SDLC Relay Standards: The SDLC Relay implementation is based on the following specifications:

- IBM SDLC General information (IBM Document GA27-3093)
- ISO 3309-1984—High-Level Data Link Control (HDLC) - Frame Structure

Bridging

The Adaptive Source Route Bridging function combines Transparent and Source Route bridging, along with translational bridging between the two. The bridging support varies from model to model, depending on the interfaces on each model (see Table 2). For a description of Source Route bridging and Source Route Transparent bridging, refer to the Distributed Routing Software V1.0A SPD (SPD 54.96.02).

Table 2
Bridging Support on Brouter Models

Product	IEEE 802.1d Transparent	Source Route	Source Route Transparent
RouteAbout Access EI	Yes		

Table 2 (Cont.)
Bridging Support on Brouter Models

Product	IEEE 802.1d Transparent	Source Route	Source Route Transparent
RouteAbout Access EW	Yes		
RouteAbout Access TW	Yes	Yes	Yes
RouteAbout Central EW	Yes		
DECswitch 900EE	Yes		
DECswitch 900EF	Yes		
DECswitch 900ET	Yes	Yes	Yes

The Bridging function interconnects networks for protocols that are not routed by the Brouter, and for protocols that do not have a Network layer to support routing. The Brouter can route some protocols, while concurrently bridging others.

Transparent Bridging: The Brouter performs Transparent bridging as defined in the IEEE 802.1d standard. It serves as a learning bridge in promiscuous mode, building a table of addresses and their source interface. These learned addresses are aged and deleted if not seen within the configurable timeout period. Permanent entries can be added to the address table. The Brouter can filter frames based on their destination address.

The Brouter participates in the IEEE 802.1d Spanning Tree Bridge algorithm, forming a loop-free network topology that connects all bridged LANs. Configuration parameters can be used to create a predictable network topology, but the default values are usually adequate. The Spanning Tree Bridge algorithm can be enabled or disabled on each port of the Brouter.

MAC Address Filtering: The RouteAbout products described in this SPD perform MAC address filtering. Frames can be filtered or tagged for Bandwidth Reservation based on the source or destination address, a mask applied to the frame, an interface number, or an input/output designation. Filters can be applied separately to each port on the Brouter. Filters can also be applied to ports established for IP tunneling. Filtered packets cannot pass through the Brouter; tagged packets are forwarded based on the configured bandwidth allocation and prioritization.

Protocol Filtering: Configuration parameters limit what protocol types are bridged in order to confine certain protocols to certain networks. These protocols can be specified based on the value in one of these fields:

- Ethernet Type
- IEEE 802.2 Destination Service Access Point
- IEEE 802 Subnetwork Access Protocol (SNAP) Protocol Identifier

IP Tunneling: Bridging across an IP network can be supported by adding a port (interface) to the Brouter bridging function. The IP tunnel provides for optimal IP routing of Transparent frames to known destinations.

When the Brouter receives a Transparent frame from the IP tunnel, it places the IP address of the source tunnel portal in the address table along with the MAC address. The Brouter can then place frames sent back to that MAC address in the correct IP portal. Similarly, the Brouter learns the Source Routing segment numbers that can be reached through each IP portal.

The IP tunnel can be configured to use IP Class D multicast via the Internet Group Multicast Protocol (IGMP) and the Multicast Extensions to Open Shortest Path First (MOSPF) routing protocol. This approach simplifies configuration, because all Routers participating in the IP tunnel need only to be configured for the same IP Class D address. IGMP and OSPF find the paths between them.

Ethernet Frame Format Translation: An

Ethernet/IEEE 802.3 network can simultaneously support use of the Ethernet and IEEE 802.3 Data Link layers. However, Ethernet frames must be translated to the IEEE 802.2 format for transmission across Token-Ring and FDDI LANs. The Brouter provides this transparency across mixed LAN types in accordance with the IEEE standards.

Bridging Standards: The bridging implementation is based on the following specifications:

- IEEE 802—LAN and MAN Overview and Architecture
- IEEE 802.1d—Media Access Control (MAC) Bridges
- RFC 1490—Multiprotocol Interconnect Over Frame Relay Networks

Frame Relay Interface

Frame Relay provides extended LAN services over a wide-area network in a point-to-point or point-to-multipoint manner. The Frame Relay interface gives access to Frame Relay services based on the Core Aspects of the LAPD data link layer protocol, ANSI T1.618-1991. The Frame Relay interface provides network addressing, congestion control, and network synchronization for Permanent Virtual Circuit (PVC) connections. Each Brouter interface supports as many as 64 PVCs.

Physical Access: Physical access is through the serial port, either point-to-point or via the network side of a CSU/DSU. Frame Relay can run up to T1/E1 speeds.

PPP Data Link Support: PPP is handled over the Frame Relay interface via an encapsulation technique. This implementation is compatible with the DECNIS implementation.

Protocol Support: The network layer protocols supported according to RFC 1490 include IP, IPX, AppleTalk, DECnet Phase IV, DECnet/OSI, and SNA. Transparent bridging is also supported as specified in RFC 1490, as is SDLC Relay encapsulated in IP. The segmentation and reassembly options are not supported.

The disable split horizon function supports AppleTalk II in partially meshed Frame Relay networks by ensuring that all routing tables are propagated from "hub routers." A hub router is a router connected to two or more routers over a partially meshed network.

The Boundary Access Node (BAN) function connects IBM Type 2.0 and Type 2.1 nodes over a Frame Relay network directly to the Frame Relay interface on the host front-end processor. Boundary Access Node eliminates the need for a router at the host site. Both BAN Type I (Bridge Mode) and BAN Type II (DLSw BAN) are supported, as detailed in the Boundary Access Node section of this SPD.

Connection Management: The Frame Relay interface supports point-to-point Permanent Virtual Circuit (PVC) connections. The state of each PVC configured on the Frame Relay interface is monitored by network management on a periodic basis for remote end point status. Connections can be configured manually, learned dynamically by network management, or both. Connections learned dynamically by network management, termed "orphan connections," can be disabled for security.

Protocol Multiplexing: Packets forwarded on the Frame Relay interface are multiplexed on different PVCs, depending on the network protocol address to PVC mapping. Network protocol addresses are mapped to PVCs dynamically by the ARP protocol. The Brouter also uses Inverse ARP to map the MAC address to the

IP, IPX, or AppleTalk address, as defined in RFC 1293. Static address assignments are also supported, allowing interoperability with Frame Relay devices that do not support ARP.

Multicast Emulation: The Frame Relay interface can be configured to emulate multicasting, allowing protocols such as ARP and RIP to function as in a LAN environment. Data Link Connection Identifiers (DLCIs) can be learned by the LMI interface or can be assigned during interface configuration. In either case, the DLCIs are treated as a list of possible logical points of attachment, all of which are candidates for directed multicast requests. The Frame Relay interface supports 2 octet DLCI addressing.

Congestion Avoidance: The Frame Relay interface avoids congestion by using Variable Information Rate (VIR) to determine the information rate for each PVC. Variable Information Rate supports three parameters per PVC: Committed Information Rate (CIR), Committed Burst Size (Bc) and Excess Burst Size (Be). The Brouter uses these parameters in conjunction with congestion monitoring to optimize PVC throughput. When the network is congested, the Brouter reduces the calculated value for each PVC's VIR to the minimum. The Brouter also determines the maximum value each PVC can transmit using the Committed Information Rate and Committed Burst Size.

Management: The Frame Relay interface complies with the Local Management Interface (LMI) protocol defined in Annex A of ITU Recommendation Q.933. LMI Interim Revision 3 as outlined in Stratacom's Frame Relay Interface specification is supported, along with the ANSI Annex D Specification.

An extension of RFC 1315 provides SNMP access to LMI information on the physical interface, as well as configuration and run time statistics on each virtual circuit on the physical interface. Traps provide information on circuit state changes.

Bandwidth Reservation Support: On the Frame Relay interface, Bandwidth Reservation can be applied separately to each virtual circuit. Bandwidth available to a circuit can be allocated to classes of traffic, protocols and filters assigned to those classes, and priority assigned to each protocol and filter in the class.

Alternatively, Bandwidth Reservation can be applied at the interface (physical DLCI) level. In this case, the bandwidth of the Frame Relay interface as a whole is allocated to classes of virtual circuits (DLCIs).

Frame Relay Standards: The Frame Relay implementation is based on the following specifications:

- ANSI T1.617—DSSI Signaling Specification for Frame Relay Bearer Service

- ANSI T1.617 Annex D—Additional Procedures for Permanent Virtual Connections (PVCs) Using Un-numbered Information Frames
- ANSI T1.618—DSSI Core Aspects of Frame Protocol for Use with Frame Relay Bearer Service
- ITU Q.933 Annex D—DSSI Signaling Specification for Frame Mode Basic Call Control
- Frame Relay Interface Specification, Revision 3.0, StrataCom, Inc., 1990
- RFC 1293—Inverse Address Resolution Protocol
- RFC 1315—Frame Relay DTE Management Information Base
- RFC 1490—Multiprotocol Interconnect Over Frame Relay Networks

X.25 Interface

The X.25 interface links Routers over public or private X.25 networks. The interface acts as a leased circuit service DTE port. The X.25 software implements the Physical, Data Link, and Network layers to transport upper layer protocol packets to other Routers. Each line can be configured separately, so that they can be connected to different X.25 networks. Each line supports as many as 227 switched virtual circuits. Priority is given to locally originated traffic.

Physical Interface: Table 3 shows the physical interfaces that are supported via the appropriate adapter cables, along with the supported baud rates.

Table 3

X.25 Physical Interfaces and Supported Baud Rates

Interface	Baud Rate	Adapter Cable
X.21	4800-256K	BC12F-06
V.35	4800-256K	BC12G-06
EIA 232/V.28	4800-19.2K	BC12L-06

Switched and Permanent Virtual Circuits: The X.25 interface initiates and manages X.25 Switched Virtual Circuits (SVCs) to transport higher level protocol data. SVCs, once opened, remain open during periods of activity, and close following a configurable period of inactivity. The X.25 interface load shares over SVCs transporting data for a higher layer protocol to a given destination DTE and attempts to open adjacent SVCs during congestion periods. Permanent Virtual Circuits (PVCs) are a configurable option. PVCs have precedence over SVCs.

Static and Dynamically Assigned Circuits: Static circuits are established when the routing circuit is enabled and remain set up until the routing circuit is disabled. Dynamically assigned circuits are established only when there is information to send over the routing circuit. They are disconnected when there is no more information to send.

Protocol Support: The X.25 interface supports the IP, IPX, and DECnet Phase IV protocols, as well as protocols such as Source Route Bridging and SDLC Relay that are encapsulated in IP.

The IP protocol can use either the encapsulation and addressing procedures specified in RFC 877 and RFC 1356, or those specified in DDN Standard X.25.

The IPX protocol is supported via use of the Call User Data (CUD) field, in accordance with RFC 1356.

The DECnet Phase IV protocol is supported via statically assigned X.25 routing circuits. DECnet Phase IV connections over X.25 are supported only to other Routers.

X.25 Standards: The X.25 implementation is based on the following specifications:

- ITU Recommendation X.25 (1980 and 1984)
- ISO 8208—X.25 Packet Level Protocol for Data Terminal Equipment
- RFC 877—Standard for the Transmission of IP Datagrams Over Public Data Networks
- RFC 1356—Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode
- Defense Data Network X.25 Host Interface Specification

Ethernet Interface

Ethernet Protocol Support: When configured for Ethernet operation, the interface provides the Physical and Data Link layers. For IEEE 802.3 operation, the interface provides the Physical and Ethernet layers. The IEEE 802.2 Logical Link Control layer is supported, but only for Class 1 (connectionless) operation.

RouteAbout Access EW and RouteAbout Access EI: The RouteAbout Access EW and the RouteAbout Access EI have an Ethernet thin wire (BNC) interface on the side and an unshielded twisted pair (RJ45) interface on the front panel. When the unit is used as a standalone Router, it detects connection to either the thin wire or the unshielded twisted pair interface and auto-configures accordingly. When installed in a DEChub 90 or a DEChub 900, the Ethernet connection is to the thin wire channel in the DEChub and the front panel interface is not used. However, the command line interface

can be used to reconfigure the unit to use the front panel Ethernet interface.

RouteAbout Central EW: The RouteAbout Central EW has two unshielded twisted pair (RJ45) interfaces on the front panel. When the unit is installed in a DEChub 900, each interface can be switched individually to a hub backplane channel. When the unit is installed in a DEChub ONE, front panel port 1 can be switched to the DEChub ONE AUI port.

Ethernet Standards: The Ethernet interface is based on the following specifications:

- Ethernet Version 2.0, Digital/Intel/Xerox, November 1982
- ANSI/IEEE 802.3—Carrier Sense Multiple Access With Collision Detection
- ISO/DIS 8802/3—Carrier Sense Multiple Access With Collision Detection
- RFC 826—Ethernet Address Resolution Protocol
- RFC 894—Standard for the Transmission of IP Datagrams over Ethernet Networks
- RFC 1293—Inverse Address Resolution Protocol
- RFC 1623—Ethernet Management Information Base

Serial Line Interface

All of the RouteAbout models described in this SPD have at least one T1/E1 serial port, as described in Table 1. This port supports line speeds from 4800 baud to 2.048 Mbs. Bandwidth Reservation is supported on serial lines.

Physical Interface: The front panel serial port connectors support the X.21, V.35, EIA232/V.28, EIA422/EIA449/V.11, and EIA530A physical interfaces, depending on the adapter cable installed (see Table 5). The Brouter automatically senses the type of cable installed and provides the appropriate physical interface.

The WAN interfaces act as DTE. NRZ and NRZI encoding are supported, with a maximum frame size of 18,000 bytes.

V.25bis Switched Circuit Support: The Brouter initiates and accepts switched circuit connections, either on demand, automatically from restart, or on command from the operator. Dial backup is also supported, as described below. Bandwidth Reservation is supported on V.25bis switched circuits.

Switched circuit support requires use of a V.25bis-compliant synchronous modem, CSU/DSU, or terminal adapter (for ISDN services). The data rate is limited by the DCE device and the carrier service used for the serial line.

Dial Backup: Routers with multiple serial lines support V.25bis for dial backup. Dial backup instructs the secondary serial line to automatically set up a dial up connection to the remote site when the primary line connection is lost. Dial backup is supported on PPP and Serial Line Protocol serial lines.

PPP Data Link: The Point-to-Point (PPP) data link is supported for the IP, IPX (IPXWAN), AppleTalk, DECnet Phase IV protocols, and DECnet/OSI protocols. Other protocols are supported if encapsulated within IP. IP header compression is supported.

The Bridging Control Protocol (BNCP) supports bridging over PPP in conformance with RFC 1220. The PPP link can be configured for Transparent bridging, Source Route bridging, or both. BNCP negotiates the source route segment number. LAT tinygram compression is supported.

The PPP Authentication Protocols (PAP) provide security through a password authentication method.

Serial Line Protocol Data Link: A proprietary implementation of the HDLC protocol, the Serial Line Protocol, transports all the supported upper layer protocols.

SDLC Data Link: The Synchronous Data Link Control (SDLC) protocol supports the DLSw, Boundary Access Node, and SDLC Relay protocols. SDLC is supported on X.21 and the other physical interfaces.

Serial Line Standards: The serial line interface is based on the following specifications:

- RFC 1220—PPP Extensions for Bridging
- RFC 1317—RS232-Like Hardware Device Management Information Base)
- RFC 1331—Point-to-Point Protocol (PPP)
- RFC 1332—PPP Internet Control Protocol (IPCP)
- RFC 1334—PPP Authentication Protocols (PAP)
- RFC 1362—Novell IPX Over Various WAN Media (IPXWAN)
- RFC 1376—PPP DECnet Phase IV Control Protocol (DNCP)
- RFC 1377—PPP OSI Network Layer Control Protocol (OSINLCP)
- RFC 1378—PPP AppleTalk Control Protocol (ATCP)
- RFC 1471—PPP Link Control Protocol Management Information Base
- RFC 1552—PPP Internetwork Packet Exchange Protocol (IPXCP)
- ITU V.25bis

ISDN Interface

The RouteAbout Access EI has a single ISDN Basic Rate interface, which provides two 64 Kbps B channels for data and a 16 Kbps D channel for signaling. The two B channels cannot be aggregated to form a higher speed link to the same destination. Bandwidth Reservation is supported on the ISDN interface.

Physical Interface: An RJ45 interface on the front panel of the RouteAbout Access EI provides an ISDN S/T interface. The S/T interface collapses the TE 1 Subscriber Terminal and the NT 2 Network Termination into a single interface. In North America, an external NT 1 is required for connection to the carrier ISDN local loop. Elsewhere in the world, the carrier provides the NT 1.

Data Link Support: The PPP and Serial Line Protocol data links are supported on the ISDN line.

Failover from the Serial Port: The RouteAbout Access EI can automatically establish an ISDN connection to the destination router when its serial line fails. In this failover mode, all multiprotocol bridging and routing is supported, just as on the serial line.

Dial on Demand for IP and OSI: Dial on Demand with an ISDN B channel as the primary circuit is supported for IP and OSI static routing. Each B channel can be routed to a separate destination.

Dial on Demand for Novell IPX: Dial on demand for Novell NetWare routing is supported as described above for IP and OSI, except that the routing need not be static. Novell server spoofing is supported, where the router at the Novell server site responds to server KeepAlive messages on behalf of the clients. This spoofing eliminates a great deal of WAN traffic.

Calling ID Verification: The RouteAbout Access EI checks the calling ID of incoming calls against an authorized list before accepting the connection.

ISDN National Standards and Switch Support: The RouteAbout Access EI has been tested for compatibility with the national standards and ISDN switches described in Table 4.

Table 4

Supported ISDN National Standards and Switches

Country	National Standard	Supported Switches
United States	National ISDN 1	AT&T 5ESS NT DMS100
Germany, United Kingdom, Belgium, Netherlands	Euro-ISDN Net3	

France	VN3
Australia	TS-013
Japan	NTT INS64

Bandwidth Reservation

Bandwidth Reservation guarantees outgoing bandwidth on serial lines, Frame Relay interfaces, V.25bis switched circuits, and ISDN lines. Both the PPP and the Serial Line Protocol data links are supported. Bandwidth Reservation is not available for X.25 interfaces. Each interface on the Brouter has its own Bandwidth Reservation settings. The system reserves percentages of the total bandwidth for specified classes of traffic. These percentages are a guaranteed minimum for the class when the serial line is fully loaded. A class can exceed its guaranteed minimum on a line with light traffic, using up to 100% of the line bandwidth. The system dynamically adapts to changes in line speed, applying the same percentage to the new line speed.

Bandwidth Classes: The serial line bandwidth can be divided among as many as 100 classes of traffic.

Priority Levels: There are four priority levels that prioritize traffic within a bandwidth class—Urgent, High, Normal, and Low. Normal is the default. The priority settings within a bandwidth class have no effect on the other classes. No bandwidth class has priority over the others.

Protocol Support: Only network protocols can be set to classes or priorities within a class, with the exception of the nine special filters discussed below. The supported protocols are IP, IPX, AppleTalk, DECnet Phase IV, and DECnet/OSI. Bridging is also supported.

MAC Address Filtering: MAC Address filtering can be used with the Bandwidth Reservation System. Frames can be filtered or tagged for Bandwidth Reservation based on the source or destination address, a mask applied to the frame, an interface number, or an input/output designation. Filters can be applied separately to each interface on the Brouter. Filtered packets cannot pass through the Brouter; tagged packets are forwarded based on the configured bandwidth allocation and prioritization.

Special Filters: The nine special filters shown below can be used as bandwidth classes or to set priority within a class. The filters can also be used to exclude traffic from the serial line. The filters are listed in order of precedence, because it is possible for a packet to be a member of several filters.

- MAC Address (Universal Filter)
- NetBIOS

- SNA
- IP Tunneling
- SDLC Relay Encapsulation (IP)
- Multicast (IP)
- SNMP (IP)
- rlogin (IP)
- telnet (IP)

Frame Relay Support: On Frame Relay interfaces, Bandwidth Reservation can be applied separately to each virtual circuit. Alternatively, Bandwidth Reservation can be applied at the interface (physical DLCI) level. In this case, the bandwidth of the Frame Relay interface as a whole is allocated to classes of virtual circuits (DLCIs).

Installation

Installation: The Distributed Routing software is factory-installed in the flash memory of the Brouter. Upgrades are performed with the TFTP protocol either locally or remotely over any supported interface. The software is provided on a CD-ROM in ISO 9660 format for DOS/Windows, OpenVMS AXP, and Digital UNIX platforms. Any system that supports IP and TFTP can serve as the load host.

For reloads when there is no valid software image in the flash memory of the Brouter, loading is supported by BOOTP/TFTP code in the Brouter PROM. The load host may be either local or remote.

In either case, routing is suspended while the software load is taking place.

Configuration and Management

The first step in configuring the Brouter is to make it remotely accessible by assigning it an IP address, subnet mask, and default gateway. For the RouteAbout Central EW, this is done by making a local terminal connection to the setup port on the DEChub ONE or DEChub 900. The setup port does not support modem control or network access of any kind. The setup port on the DEChub 900 provides a menu from the hub Management Access Module (MAM), through which you can attach to the module and see the RouteAbout Central EW's own setup menu. The setup port on the DEChub ONE takes the user directly to the RouteAbout Central EW setup menu. The RouteAbout Central EW setup menu minimally configures the Brouter so that it can send and receive IP datagrams.

For the RouteAbout Access EW and the RouteAbout Access EI, a direct connection to the Brouter console port is required regardless of whether or not it is installed in a hub. A set up menu minimally configures the Brouter so that it can send and receive IP datagrams. The user is then transferred into the Distributed Routing software command line interface (CLI).

For all of the products, the remainder of the Brouter configuration is performed using the Distributed Routing software command line interface.

The Distributed Routing software command line interface (CLI) can be launched from HUBwatch, accessed remotely via TCP/IP TELNET, or accessed locally through an EIA 232 terminal connected to the RJ45 console port of the RouteAbout Access EW or the RouteAbout Access EI. Dial in access to the RouteAbout Access EW or EI console port via a modem is also supported. The RouteAbout Central EW does not support console port access. Telnet Client allows users to access Telnet servers from the Brouter console CLI. As many as three outbound Telnet sessions are supported.

The following statistics are available:

- Uptime
- Last restart or reboot
- Memory utilization
- Packets forwarded
- Detailed interface error counters
- Protocol status and error counters

Quickconfig Utility: The Quickconfig utility allows for fast configuration of the Brouter interfaces, IP addresses, bridging, and IP, IPX, and DECnet Phase IV protocol support.

HUBwatch: HUBwatch selects the hub channel to which the Brouter LAN backplane interface is attached. HUBwatch also displays an icon of the Brouter within a display of the hub slot configuration. By clicking on the icon, the user can launch the Distributed Routing software command line interface. Configuration and monitoring using HUBwatch can be done over the network or via a TCP/IP SLIP connection to the DEChub Out of Band Management (OBM) Port.

SNMP: SNMP provides a method of monitoring and managing the operation of the Brouter remotely, using a standardized, extensible UDP-based protocol. It can examine the state of the Brouter, collect various statistics, and generate trap messages. The complete MIB-II is provided with the exception of ifInNUcastPkts, ifOutNUcastPkts, and the TCP group. The Address Translation and Routing tables are not settable.

SNMP Standards: The SNMP implementation is based on the following RFCs. SNMP Gets and Traps are supported for the MIBs listed below:

- RFC 1155—Structure and Identification of Management Information for TCP/IP-Based Internets
- RFC 1157—Simple Network Management Protocol (SNMP)
- RFC 1213—Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II
- RFC 1231—Token-Ring MIB
- RFC 1243—AppleTalk MIB
- RFC 1253—OSPF Version 2 MIB
- RFC 1284—Ethernet MIB
- RFC 1285—FDDI MIB
- RFC 1286—Bridge MIB
- RFC 1315—Frame Relay DTE MIB
- RFC 1317—RS232-Like Hardware Device MIB
- RFC 1471—PPP Link Control Protocol MIB
- RFC 1573—Evolution of the Interfaces Group of MIB-II
- Distributed Routing Software MIB (Resource Group, ELS Group, and Traps)
- DLSw MIB (IBM Document GC30-3567)
- Novell IPX MIB
- Novell RIP-SAP MIB

Error statistics are collected on a node and port basis for both Token Ring and Ethernet. In addition, the ring ordered node list, otherwise known as NUAN (Nearest Upstream Active Node) list, is maintained. System delta events are captured in an event log for retrieval by a SNMP network management system. Automatic mapping of MAC address to physical port is provided for Token Ring only. For Token Ring, certain events such as ring beaconing cause the SNMP agent to automatically isolate the faulty port and remove it from the ring.

For security, specific portions of the MIB or the entire MIB can be assigned to a community. Each community has a list of IP addresses that can access the community and/or receive traps.

INSTALLATION

If the Brouter is to be connected to a public X.25 network, Digital recommends that a customer's first purchase of the product include Digital Installation Services. These services provide installation of the software by an experienced Digital software specialist.

HARDWARE REQUIREMENTS

The Distributed Routing software requires the following hardware:

Brouter Hardware Unit: A RouteAbout Access EW, RouteAbout Access EI, or RouteAbout Central EW hardware unit is required.

DEChub: DEChub variants of the RouteAbout Access EW and RouteAbout Access EI require a DEChub 90 or a DEChub 900. Stand-Alone and MultiStack variants do not require additional hardware, as they are supplied with a power supply and power cord.

The RouteAbout Central EW requires either a DEChub 900 for hub configuration or a DEChub ONE for stand-alone use.

Console Terminal: A terminal is required for local configuration of the Brouter.

Ethernet Cables: For the Ethernet interface on a RouteAbout Access EW or the RouteAbout Access EI, a 10-Base-2 (BNC) or 10-Base-T (RJ45) Ethernet connection is required.

On the RouteAbout Central EW, a 10-Base-T (RJ45) Ethernet connection is required for each Ethernet interface on the front panel.

Serial Port Cables: An adapter cable is required for each serial interface on the Brouter. This cable provides the required physical interface, as shown in Table 5.

Table 5
Physical Interfaces and Adapter Cables

Interface	Adapter Cable
X.21	BC12F-06
V.35	BC12G-06
EIA 232/V.28	BC12L-06
EIA 530A	BC12J-06
EIA 422/EIA 449/V.11	BC12H-06

DCE Device for Serial Lines: A DCE device (DSU/CSU or modem) is required for each synchronous line. Switched circuit support requires use of a V.25bis-compliant synchronous modem, CSU/DSU, or terminal adapter (for ISDN services).

NT 1 for ISDN Lines: In North America, an NT 1 Network Termination is required for connecting the RouteAbout access EI ISDN line to the ISDN carrier local loop. Elsewhere in the world, the carrier provides the NT 1.

SOFTWARE REQUIREMENTS

A TCP/IP System is required for remote configuration and management of the Brouter via TELNET. TFTP support is required for performing reloads and upgrades of the Brouter software and for receiving dumps.

OPTIONAL SOFTWARE

An SNMP Network Management Station is required for monitoring of SNMP variables.

A HUBwatch system is required for configuring the hub backplane channel attachment of certain modules in DEChub configurations. The RouteAbout Access EW requires HUBwatch Version 4.1 or later. The RouteAbout Access EI and RouteAbout Central EW require

HUBwatch Version 5.0 or later.

GROWTH CONSIDERATIONS

The minimum hardware/software requirements for any future version of this product may be different from the requirements for the current version.

DISTRIBUTION MEDIA

The Brouter software is distributed on a CD-ROM in ISO 9660 format for DOS/Windows, OpenVMS AXP, and Digital UNIX platforms. The product documentation is also provided on a CD-ROM, along with a *DynaText* reader. Table 6 describes the system requirements for the *DynaText* reader.

Table 6
***DynaText* System Requirements**

Platform	Operating System	Windowing System	RAM	Disk Space
IBM PC or compatible; 386sx or greater	MS-DOS 5.0 or greater	MS Windows 3.1 in enhanced mode	4MB	4MB
Sun SPARC system	SunOS 4.1.x or Solaris 2.2	X11R4, X11R5, Motif 1.2.x	12MB	4MB
IBM RS/6000	AIX 3.2.5 or later	X11R4, X11R5, Motif 1.2.x	12MB	3MB
HP9000 Series 700	HPUX 9.0 or later	X11R4, X11R5, Motif 1.2.x	12MB	3MB
Digital Alpha	Digital UNIX V3.0 or later ¹	X11R4, X11R5, Motif 1.1.x	12MB	3MB

¹Requires that you install the C++ Run-Time Library option.

ORDERING INFORMATION

For initial purchase of a Brouter, use the DE***-** option number shown in Table 7. This part number is a complete package including the hardware unit with the software preloaded in the flash memory and the software license. CD-ROMs shipped with each Brouter contain the software, the software documentation, and a *DynaText* documentation reader. A hardware installation manual and a set of software quick reference cards are also included. To purchase a complete hardcopy documentation set, use the GZ kit part number listed below.

Table 7
Distributed Routing Software V1.1 Ordering Information

	RouteAbout Access EW	RouteAbout Access EI	RouteAbout Central EW
Initial Ordering			
Brouter Unit with IP Software			
Stand-Alone	DEX2R-F*	DEXBR-F*	N/A
Hub	DEX2R-MB	DEXBR-MB	DEZ8R-S*
MultiStack	DEX2R-T*	DEXBR-T*	N/A
Brouter Unit with MP Software			
Stand-Alone	DEX2R-D*	DEXBR-D*	N/A
Hub	DEX2R-MA	DEXBR-MA	DEZ8R-R*
MultiStack	DEX2R-S*	DEXBR-S*	N/A
Updates and Upgrades			
IP Software Update License	QL-4C1A9-RA	QL-4SQA9-RA	QL-4P2A9-RA
Multiprotocol Software Update License	QL-4C2A9-RA	QL-4SRA9-RA	QL-4P3A9-RA
IP to Multiprotocol Upgrade License	QL-4C3A9-AA	QL-4SSA9-AA	QL-4P4A9-AA
IP Software Media and Documentation	QA-4C1AA-H8	QA-4SQAA-H8	QA-4P2AA-H8
MP Software Media and Documentation	QA-4C2AA-H8	QA-4SRAA-H8	QA-4P3AA-H8
Hardcopy Documentation and Services			
Hardcopy Software Documentation	QA-4P3AA-GZ	QA-4P3AA-GZ	QA-4P3AA-GZ
IP Software Product Services	QT-4C1A*.**	QT-4SQA*.**	QT-4P2A*.**
MP Software Product Services	QT-4C2A*.**	QT-4SRA*.**	QT-4P3A*.**

*Denotes the country kit variant or memory option. For additional information on available country kits, memory options, licenses, services, and media, refer to the appropriate price book.

IP Software License: The IP Software license provides IP, bridging, and all WAN services including PPP, Frame Relay, X.25, V.25bis, and Bandwidth Reservation.

Multiprotocol Software License: The Multiprotocol Software license includes all the IP Software protocols, as well as Novell IPX, AppleTalk, DECnet Phase IV, DECnet/OSI, DLSw, and SDLC Relay.

Upgrade License: An IP to Multiprotocol Software Upgrade license is available for all the products. Customers who are upgrading should also purchase the Multiprotocol Software Media and Documentation kit for the Brouter.

SOFTWARE LICENSING

A separate license is required for each Brouter hardware unit on which the software product is to be used. This license is included in the price of the Brouter hardware. A license letter shipped with the hardware unit,

along with the invoice, serves as proof of license. The software license may also be purchased separately.

The licensing provisions of Digital's Standard Terms and Conditions specify that the software and any part thereof (but excluding those parts specific to the load hosts) may be used only on the single Brouter hardware unit on which the software is operated, but may be copied, in whole or in part (with the proper inclusion of Digital's copy right notice and any proprietary notices on the software) between multiple load hosts on the same LAN.

This software is furnished under the licensing provisions of Digital's Standard Terms and Conditions. For more information about Digital's licensing terms and policies, contact your local Digital office.

SOFTWARE PRODUCT SERVICES

A variety of service options are available from Digital. For more information, contact your local Digital office.

SOFTWARE WARRANTY

Warranty for this software product is provided by Digital with the purchase of a license for the product as defined in the Software Warranty Addendum of this SPD.

The above information is valid at time of release. Please contact your local Digital office for the most up-to-date information.

- ® AIX, IBM, and NetView are registered trademarks of International Business Machines Corporation.
- ® AppleTalk, EtherTalk, and TokenTalk are registered trademarks of Apple Computer, Inc.
- ® DOS, MS-DOS, and Windows are registered trademarks of Microsoft Corporation.
- ® DynaText is a registered trademark of Electronic Book Technologies, Inc.
- ® HP is a registered trademark of the Hewlett-Packard Company.
- ® Motif is a registered trademark of the Open Software Foundation, Inc.
- ® NetWare and Novell are registered trademarks of Novell, Inc.
- ® Solaris and Sun are registered trademarks and SunOS is a trademark of Sun Microsystems, Inc.
- ® UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.
- ™ Alpha, AXP, DECCONNECT, DEChub, DEChub ONE, DECnet, DECswitch, Digital, HUBwatch, MultiStack, OpenVMS, RouteAbout, VMS, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

© 1996 Digital Equipment Corporation. All rights reserved.