

Routing and Networking Overview

Order No. AA-HS15A-TK

December 1986

The *Routing and Networking Overview* is an introduction to routing in DECnet networks, written for network managers and others responsible for configuring and maintaining networks. It defines the basic routing terminology and concepts. This manual also describes the types of computer systems that perform routing, explains how they carry out their routing functions, and provides guidelines for configuring networks for optimal routing performance.

Supersession/Update Information: This is a new manual.

Software Version: DECrouter 200 V1.0

digital

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may only be used or copied in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by Digital or its affiliated companies.

Copyright © 1986 by Digital Equipment Corporation
All Rights Reserved.
Printed in U.S.A.

The postage-prepaid Readers Comments form on the last page of this document requests the user's critical evaluation to assist us in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

DEC	Micro/R SX	TOPS-10
DECconnect	MicroVAX	TOPS-20
DECmate	MicroVMS	ULTRIX-32
DECnet	PDP	ULTRIX-32m
DECserver	P/OS	UNIBUS
DECtape	Professional	VAX
DECUS	Rainbow	VAXcluster
DECwriter	RSTS	VAXmate
DIBOL	RSX	VAX/VMS
digital	RSX-11M-PLUS	VMS
MASSBUS	RT	VT
MicroPDP-11	ThinWire	Work Processor

Bell is a trademark of Bell Telephone Companies.

IBM is a registered trademark of International Business Machines Corporation.

PC/XT and Personal Computer AT are trademarks of International Business Machines Corporation.

TEFLON is a trademark of Dupont.

This manual was produced by Networks and Communications Publications.

Contents

Preface

1 Routing: Introduction

1.1	What Is Routing?	1-1
1.2	How Routers Contribute to Networking	1-2
1.3	Host-Based Routers and Server-Based Routers	1-5

2 DECnet Routing and Enhancements

2.1	Adaptive Routing	2-1
2.2	Routing Terms	2-2
2.3	Routing Features	2-6
2.4	Routing Control Messages	2-8
2.4.1	Segmented Routing Messages	2-9
2.4.2	Timing of Routing Message Transmissions	2-9
2.5	Routing Parameters	2-9
2.5.1	Circuit Cost Parameter	2-10
2.5.2	Maximum Path-Control Parameters	2-12
2.5.3	Route-Through Control Parameter	2-13
2.6	DECnet Phase IV Routing Enhancements	2-13
2.6.1	Area Routing Improvements	2-13
2.6.2	Path Splitting	2-13

3 Network Technologies and Topologies: Ways of Building Networks

3.1	DECnet Environments	3-3
3.1.1	Wide Area Networks	3-3
3.1.2	Local Area Networks	3-3
3.2	Data Link Technologies	3-5
3.2.1	Ethernet Connections	3-6
3.2.1.1	Access Control	3-6

3.2.1.2	How Routing Works on an Ethernet LAN	3-6
3.2.2	DDCMP Connections	3-7
3.3	Multiple Area Networks	3-10
3.4	Node Characteristics	3-11
3.4.1	DECnet Phase III and Phase IV Nodes	3-11
3.4.2	Routing Capabilities	3-12
3.5	Routers and Other Products That Extend Ethernet LANs.....	3-12
3.5.1	Bridges	3-14
3.5.2	Routers	3-15
3.5.3	When to Use Bridges and Router Servers.....	3-16

4 Configuring Network Nodes for Optimal Routing Performance

4.1	General Guidelines for Configuring Networks	4-2
4.1.1	Specifying Node Identification	4-3
4.1.2	Assigning the Maximum Node Address on Nodes	4-3
4.1.3	Selecting the Buffer Size and Segment Buffer Size on Nodes	4-4
4.1.4	Assigning Circuit Costs	4-5
4.1.5	Choosing Path Lengths	4-6
4.1.5.1	Defining the MAXIMUM HOPS and MAXIMUM COST Parameters	4-7
4.1.5.2	Defining the AREA MAXIMUM COST and AREA MAXIMUM HOPS Parameters	4-8
4.1.6	Distributing Network Traffic	4-8
4.2	Guidelines for Configuring Ethernet LANs	4-9
4.2.1	Number of End Nodes (Nonrouters) on an Ethernet LAN.....	4-9
4.2.2	Number of Routers on an Ethernet LAN	4-10
4.2.3	Potential Designated Routers	4-13
4.3	Guidelines for Configuring Multiple Area Networks	4-13
4.4	When and Where to Include Routers	4-15
4.4.1	Ethernet LANs	4-15
4.4.2	Multiple Area Networks	4-16

5 Routing-related Problems in Networks

5.1	General Routing Problems	5-1
5.1.1	Connectivity Problems	5-1
5.1.2	Data Link Problems	5-5
5.1.3	Performance or Lost Packet Problems.....	5-6
5.2	Problems in Multiple Area Networks	5-10

A Configuring a Multiple Area Network

A.1	Converting a Standard Network into a Multiple Area Network	A-1
A.2	Designing Multiple Area Networks	A-4

Glossary

Figures

1-1	A Network Without Routing Nodes	1-3
1-2	A Network with Routing Nodes	1-4
1-3	Routing Offers Reliability and Flexibility	1-5
1-4	The Server-Based Router Off-loads Routing Functions from Other Nodes	1-7
2-1	Routing Paths	2-2
2-2	Routing Terms	2-3
2-3	A Multiple Area Network	2-5
2-4	Network Circuit Costs and Path Hops	2-11
2-5	Path Splitting	2-15
3-1	Synchronous Routers	3-9
3-2	Asynchronous Routers	3-10
3-3	Functional Scopes of Repeaters, Bridges, and Routers	3-14
4-1	Low-Hop vs Multiple-Hop Paths	4-7
4-2	Defining the MAXIMUM BROADCAST ROUTER Parameter	4-12
5-1	Nodes with Equal Addresses	5-2
5-2	Area Leakage	5-11
5-3	Improper Router Configuration Within an Area	5-13
5-4	Improper Configuration of Phase III and Phase IV Nodes	5-15

Tables

3-1	Comparison of Routers and Bridges	3-13
-----	---	------



Preface

The *Routing and Networking Overview* is an introduction to routing in DECnet networks. It defines basic routing terminology and concepts. This manual also describes the types of computer systems that perform routing, explains how they carry out their routing functions, and provides guidelines for configuring networks for optimal routing performance.

Intended Audience

This manual contains information intended for network managers and others responsible for configuring and maintaining networks. The manual assumes you understand and have some experience of:

- Local Area Networks (LANs)
- Wide Area Networks (WANs)
- DECnet Phase IV

Structure of the Manual

- Chapter 1 Introduces basic routing terminology and discusses the important role that routing plays in networks. Distinguishes various types of routers.
- Chapter 2 Defines basic DECnet routing terminology and concepts. Discusses the routing features provided by DECnet, including enhancements to DECnet Phase IV routing. Also discusses several DECnet routing parameters.
- Chapter 3 Describes various technologies and topologies that can combine to build a network. Compares routers, bridges, and other products for extending LANs.

- Chapter 4 Gives guidelines for configuring networks and optimizing routing performance.
- Chapter 5 Explains how to isolate problems that may develop in networks with routing, and how to avoid and solve problems.
- Appendix A Contains guidelines for configuring multiple area networks.
- Glossary Defines common routing terms and other networking terms.

Related Documents

Introduction to DECnet Phase IV

Digital's Networks: An Architecture with a Future

Overview of Digital Networking Products

VAX/VMS Networking Manual

DECnet-RSX Network Management Concepts and Procedures

Also refer to the latest edition of the *Networks and Communications Buyer's Guide*. This book is updated several times each year.

Conventions Used in This Manual

Convention	Meaning
<i>lowercase italic</i>	Variables in command lines are printed in lowercase italic type. The value is specified by the user.
boldface	Used for emphasis. Terms in boldface are included in the glossary.
UPPERCASE	Indicates name of a parameter. In command lines, indicates characters to be typed literally.

Routing plays an important part in the flow of data through the network. As network manager, you are responsible for overseeing the operation of the network as a whole. You must ensure that routing nodes are configured strategically and that routing control parameters provide for efficient data flow through the network. This chapter defines routing and routing nodes and describes the importance of routing.

1.1 What Is Routing?

Routing is the network function that determines the path or route along which data (called “packets” in this context) travels to its destination. The routing layer in the **Digital Network Architecture (DNA)** corresponds to the network layer in the **International Standards Organization** architecture (ISO).

Within a network, nodes are joined together by physical communication links over which data passes from one node to the next. These are usually referred to as **lines**, but may involve satellite, microwave, or fiber-optic links. Nodes with a direct physical link between them (called **adjacent nodes**) can communicate directly. Nonadjacent nodes can communicate only if intervening nodes forward the data along the path between the source and the destination. Intervening nodes that receive and forward data from one node to another are known as **routing nodes**.

Routing nodes have circuits that connect to two or more other nodes in the network. A **circuit** is a routing-level (logical) communications link between two adjacent nodes connected by a physical line or cable. DECnet routing nodes include network software which maintains a database describing the availability of paths to destination nodes and performs the routing function. In contrast, nonrouting nodes, also called **end nodes**, are connected to only one other node. They do not include network software for routing and cannot receive and forward data intended for other nodes.

Networks such as DECnet allow messages to be sent between any two nodes in the network, even if they are not directly connected to one another. DECnet provides **adaptive routing**, which means messages are routed through the network along the most cost-effective path currently available, and messages are rerouted automatically if a circuit becomes disabled.

A routing node's choice of path depends on the current information stored in its routing database. Routing nodes keep their routing databases up to date by regularly exchanging routing information, such as the **cost** and the number of **hops** (the number of nodes along the path to each destination) involved in sending packets to other nodes. The packets are routed to the destination along the path with the least cost. Chapter 2 defines these and other basic routing terms and discusses in further detail how DECnet routing works.

1.2 How Routers Contribute to Networking

Routing is similar to the telephone system. Messages pass from a source node (where you make the call) to a destination node (the number you are calling). Due to the complexity of the system and the fact that you can call almost any number in the world, you do not need to know the route taken by your call. To reach the number you are calling, the call is routed through several exchanges which pass the call from one area to another. Depending on the availability of lines, the route may be different each time you call the number.

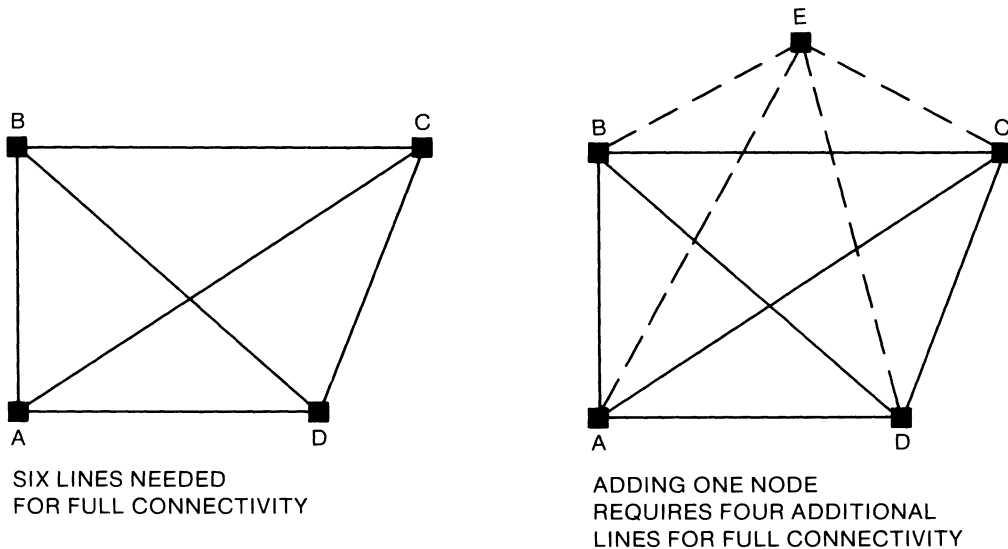
Some features that routing gives to network users are:

- Lower configuration costs
- More efficient use of network components, such as nodes and lines
- Increased reliability
- Greater flexibility for internode communications
- Congestion control

Without routing, a point-to-point network would have very high configuration costs because of the additional communication lines and hardware required. Furthermore, to add a node that is to communicate with every other node in the network, you would have to increase the number of communication lines exponentially.

Figure 1-1 shows such a point-to-point network consisting of four nodes. It requires six lines to provide full **connectivity** (where each node can communicate with every other node). Considering the communications devices needed at both ends of each line, that means up to 12 point-to-point devices can be needed. To add a fifth node (node E) and maintain full connectivity, up to four lines and eight devices would have to be added.

Note that this discussion pertains to point-to-point wide area networks (WANs), not to local area networks (LANs) such as Ethernet. An Ethernet can provide full connectivity to all nodes connected to it, and all those nodes can be configured as nonrouting nodes. Also, new nodes can be added without the exponential increase in devices. Other multi-access devices, such as multipoint lines, also share this latter feature. The point-to-point, multipoint, and Ethernet technologies are discussed further in Section 3.2.



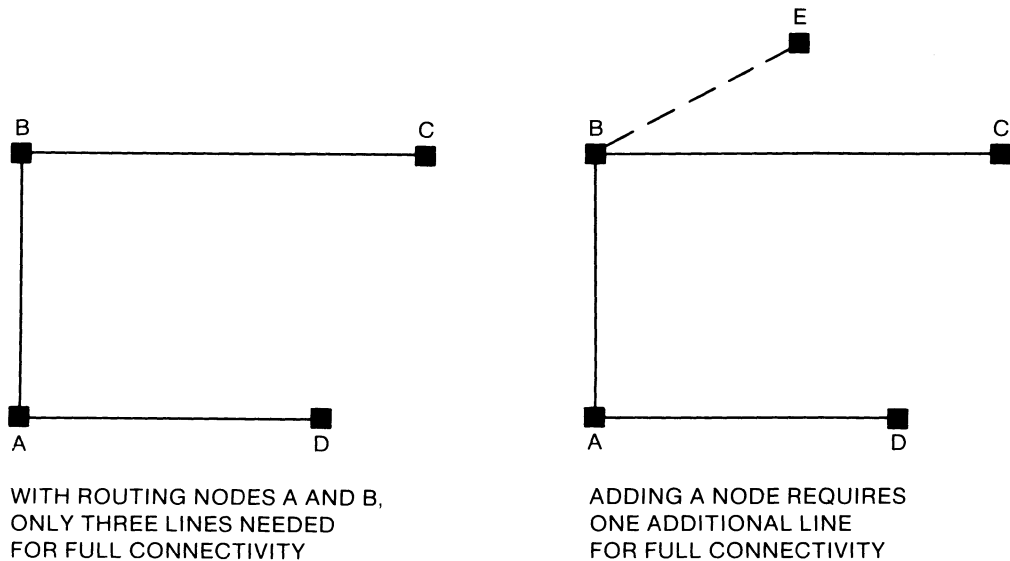
LKG-0554

Figure 1-1: A Network Without Routing Nodes

Routing nodes add flexibility to networks because all nodes do not have to be connected directly to each other. For example, consider again the four nodes A, B, C, and D. As shown in Figure 1-2, by configuring nodes A and B as routing nodes, only three lines are needed to provide full connectivity for the four nodes. In this configuration, no physical lines are needed between B and D, A and C, and C and D. One or more of the physical lines used to connect the nodes can be shared for communications between users of the network.

For instance, the line between nodes A and B enables communication between D and B, D and C, and A and B. Thus, as a shared resource, the line is more cost-effective. In contrast, referring back to Figure 1-1, the line between nodes A and B is limited to communication between A and B only, a less efficient use of the line.


Figure 1-2 also shows the minimal impact of adding a node to the network. If you connect the new node (node E) to one of the routing nodes, only one new line is needed. Node E can still communicate with every other node in the network.



LKG-0555

Figure 1-2: A Network with Routing Nodes

Another advantage of routing is that it can provide more network reliability. In the routing network shown in Figure 1-3, alternate paths are available for communications between any pair of nodes. For example, node A can communicate with node C over the path that includes node D. If that path becomes unavailable (if line AD or line DC is unavailable, or if node D fails), communications can still continue: node A can communicate with node C in the other direction, using the path that includes node B.



EACH NODE HAS
AN ALTERNATIVE PATH
TO TAKE TO ANY
OTHER NODE

LKG-0556

Figure 1-3: Routing Offers Reliability and Flexibility

1.3 Host-Based Routers and Server-Based Routers

Two types of routing nodes are host-based routers and server-based routers. **Host-based routers** are full-function nodes that handle user applications as well as routing. **Server-based routers** are nodes dedicated to routing only. Server-based routers are designed for the single function of routing.

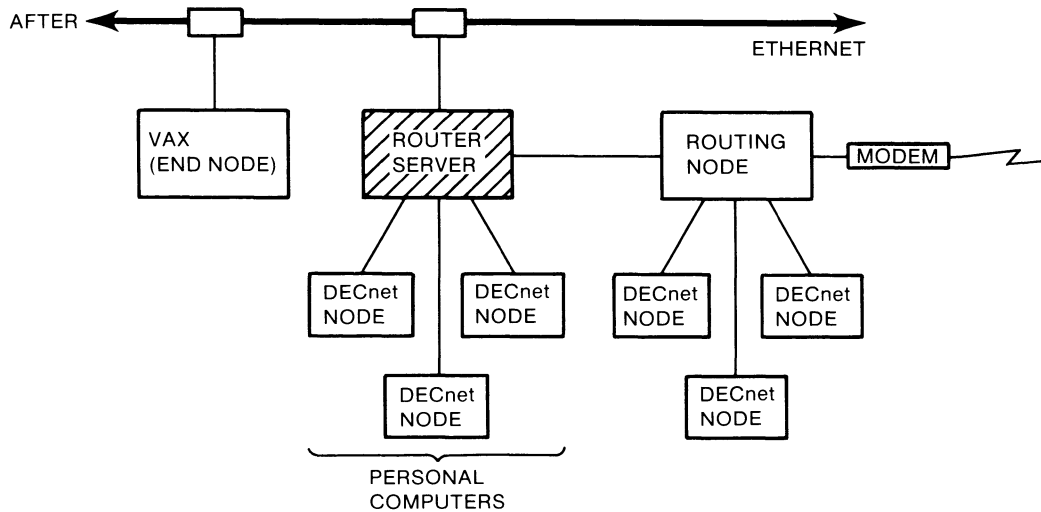
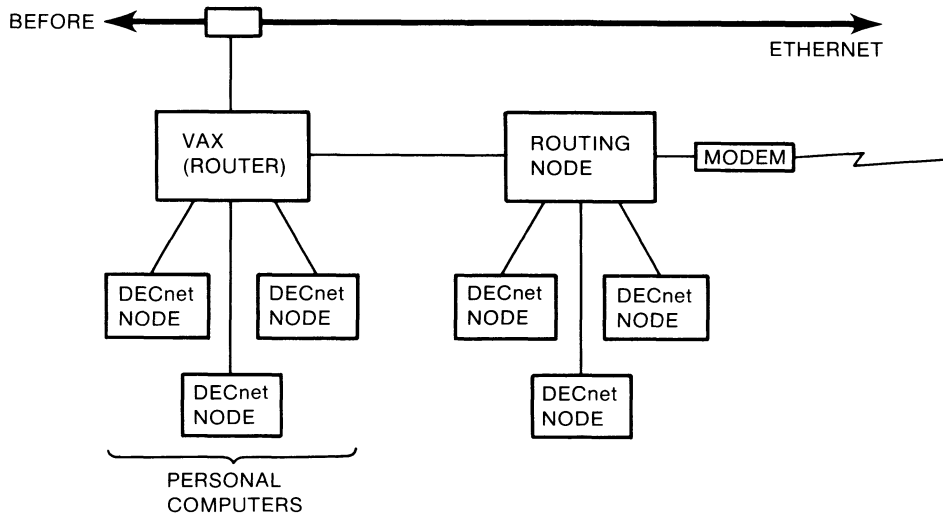
On host-based routers, user applications must compete with the routing function for CPU time and memory resources. The routing function uses a proportion of the CPU cycles, and therefore, reduces the number of cycles available for any other applications on the node. (The exact proportion used for routing depends on network traffic and the number of network connections.) Conversely, the applications may reduce the number of cycles available for the routing function. Thus, both user applications and routing may suffer when user activity and network traffic are excessive.

Server-based routers solve this problem by off-loading the routing function from the full-function nodes of the network. Figure 1-4 shows how a server-based router contributes to an Ethernet network. Without the server-based router, the VAX node in the original configuration shown in the figure is loaded with routing responsibilities for the four nodes connected to it. Notice that the routing node connected to the VAX probably adds a great deal to the VAX system's routing responsibilities, since there is likely to be heavy traffic between nodes on the Ethernet and nodes connected to the routing node.

As the figure shows, however, connecting a server-based router to the Ethernet makes routing more cost-effective than when performed by a full-function (general purpose) computer system. This is because the server-based router is designed specifically for routing. Also, the server-based router significantly reduces the load on the VAX node. All nodes previously connected to the VAX can now be connected to the server-based router.

By consolidating the routing function, the server-based router lets you configure the VAX (and any other nodes on the Ethernet) as an end node. Because end nodes do not require routing software, you save costs. Therefore, you can afford to build larger networks by including a proportionately larger number of end nodes. Also, because end nodes are free from routing responsibilities, they can dedicate more of their resources to user applications.





LKG-0557

Figure 1-4: The Server-Based Router Off-loads Routing Functions from Other Nodes




2 DECnet Routing and Enhancements

This chapter discusses DECnet routing terminology and concepts and DECnet Phase IV enhancements.

2.1 Adaptive Routing

DECnet adapts to changing conditions in the network, a capability known as **adaptive routing**. To route data from one node to another, it selects the best path currently available to the destination. In many cases, multiple paths exist between network nodes, as shown in Figure 2-1. When the primary path is unavailable, DECnet can detect the situation and redirect the data over the next best alternative path.



LKG-0558


Figure 2-1: Routing Paths

With adaptive routing, network users need not be concerned about the status of paths to a destination. Users need only specify the name of the remote node with which they wish to communicate. The DECnet routing mechanism does the rest. Each routing node determines the best path to the destination according to up-to-date information kept in a routing database on each routing node.

Routing nodes keep their routing databases up to date by regularly exchanging routing information. Whenever this routing information changes (for instance, when a circuit fails), new routing messages will be sent automatically. If a line fails, nodes affected by that line recompute their routing information and redirect data traffic over the best alternative path.

2.2 Routing Terms


The routing database kept on each routing node includes the distance and the cost to each remote node. The total distance between a source node and a destination node is called the **path length**. The path length is measured in **hops**. In Figure 2-2, the distance between node A and node B is one hop. The shortest path from node A to node D is two hops (path A-B-D). Figure 2-2 and the accompanying table also show two other paths between node A and node D.



Legend:



= node



= circuit



= circuit cost



= hop

Node A wants to send a packet to Node D. There are three possible paths.		
PATH	PATH COST	PATH LENGTH
(A) to (B), (B) to (C), (C) to (D)	$2 + 2 + 3 = 7^*$	3 hops
(A) to (B), (B) to (D)	$2 + 7 = 9$	2 hops
(A) to (B), (B) to (F), (F) to (E), (E) to (D)	$2 + 3 + 4 + 2 = 11$	4 hops

*7 is the lowest path cost; Node A therefore routes the packet to Node D via this path.

LKG-0559

Figure 2-2: Routing Terms

To optimize performance, the maximum number of hops along a path is limited to 30 by DECnet Phase IV. A lower maximum may be imposed by the network manager to prevent the use of undesirably long paths. The parameter that defines the maximum path length in the network is called the MAXIMUM HOPS parameter.

Cost is an arbitrary integer used to control the path of packets. The system manager of a node assigns a cost to each circuit of the node when generating the node's network database. The system manager usually confers with the network manager before assigning costs.

The cost of each path (**path cost**) between a source node and a destination node is the sum of the costs assigned to the circuits that compose the path. The routing function routes packets on the path of least cost, even if it is not the path with the fewest hops. Referring again to Figure 2-2, the least path cost from node A to node D is 7, along path A-B-C-D.

If two paths to a destination have the same least cost, then the routing path is selected arbitrarily. (If path splitting is enabled, a routing enhancement to DECnet Phase IV discussed in Section 2.6.2, then the routing function transmits data to the destination by dividing it evenly over the two or more paths having equal path costs.)

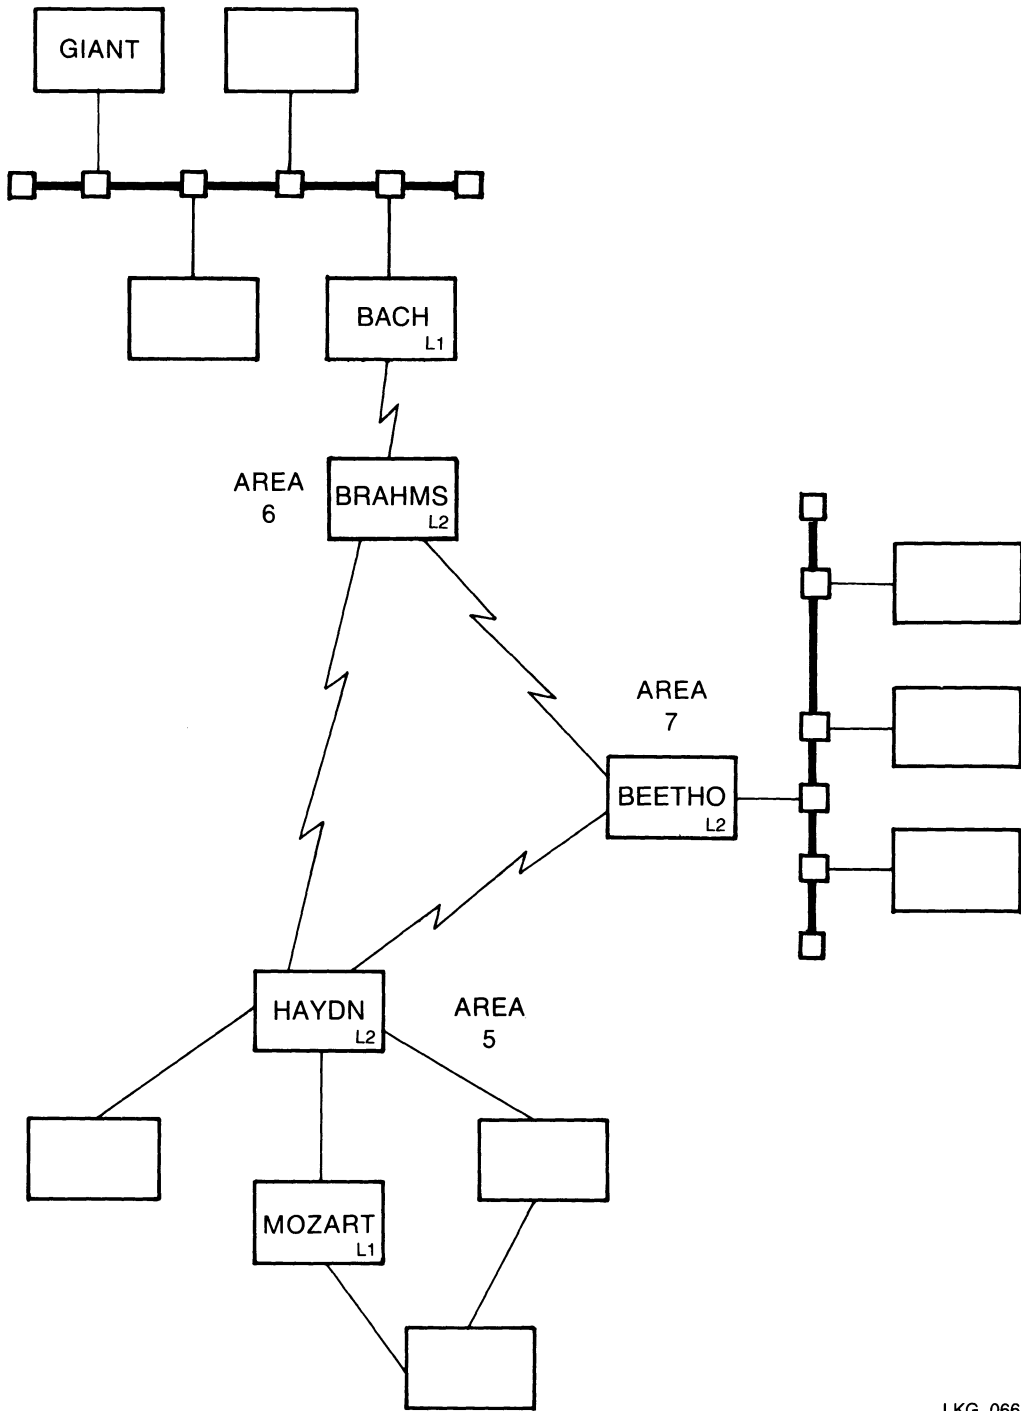
By changing the cost of a circuit, the network manager can control the flow of data through the network. For example, if traffic on one path to a destination is heavy, causing **congestion**, the network manager can increase circuit costs along that path to divert traffic to another path that is being used less. As another example, circuit costs can be changed temporarily to divert traffic for the benefit of a specific network job.

Sections 2.5.1 and 2.5.2 discuss path costs and path lengths in further detail.

The design of DECnet's routing layer gives network managers great flexibility in configuring networks. Node locations and the assignment of circuit costs and maximum-hop values are left entirely to a network manager's discretion. Thus, a DECnet network can be tailored to meet the varying needs of differing application requirements.

DECnet Phase IV supports **area routing**. Area routing permits configuration of very large networks. A large network can be divided into 63 separate **areas**, each containing up to 1023 nodes. Areas are connected by **area routers**, also called **level 2 routers**. **Level 1 routers** perform the standard routing between nodes of the same area.

Figure 2-3 shows a multiple area network. (An "L1" in figures indicates a level 1 router. An "L2" indicates a level 2 router.)



LKG-0667

Figure 2-3: A Multiple Area Network

By using areas, routing traffic within a large network is reduced. Inter-area communication is restricted to level 2 routers. Level 1 routers only perform routing to, and exchange information about, nodes in their own area. (A level 2 router also performs level 1 routing, exchanging information about nodes within its own area.)

When a node in one area wishes to communicate with a node in another area, a level 1 router in the local area forwards the request to the nearest (least path cost) level 2 router in the local area. The level 2 router keeps track of the best paths to each remote area of the network.

Referring to Figure 2-3, if node BACH in area 6 wants to send data to node MOZART in area 5, it sends the data to the nearest level 2 router in area 6, node BRAHMS. In turn, node BRAHMS uses level 2 routing to forward the data on the best path to area 5, which terminates at the level 2 router, node HAYDN. Then node HAYDN routes the data by level 1 routing to node MOZART. Notice that the level 2 routers BEETHO, BRAHMS, and HAYDN form a subnetwork to link the three areas shown in the figure.

In multiple area networks, routing parameters specific to area routing can be assigned when installing level 2 routers. Area routing parameters include:

- **MAXIMUM AREA**
Specifies the largest area number in the network (see Section 4.1.2).
- **AREA MAXIMUM HOPS**
Specifies the maximum path length between areas (see Section 2.5.2).
- **AREA MAXIMUM COST**
Specifies the maximum path cost between areas (see Section 2.5.2).

2.3 Routing Features

In addition to determining packet paths, path length, and path cost, the DECnet routing mechanism performs other functions that ensure efficient data transfer over the network. The following list summarizes all the major features:

- **Forwarding packets** — If a packet is addressed to the local node, the routing function on that node delivers it to the end-to-end communications layer. If a packet is addressed to a remote node, routing forwards it to the next adjacent node on the best available path. In multiple area networks, level 1 routers forward packets between nodes within an area; level 2 routers forward packets between areas.
- **Adapting to changes in network topology** — If a circuit or node in a path fails, routing finds an available alternative path (if there is one).
- **Adapting to different kinds of circuits** — Routing operates over paths consisting of multiple types of data-link connections, including DDCMP point-to-point and multipoint lines, Ethernet links, and X.25 virtual circuits.

- **Periodically updating routing databases on other nodes** — Routing databases on other nodes are periodically updated to reflect any topology changes; for example, circuits going on or off, and operators modifying routing parameters (see above).

- **Returning packets addressed to unreachable nodes** — If requested, routing returns packets addressed to unreachable nodes to the end-to-end communications layer, thereby keeping it informed of the status of packet transmissions (otherwise, that layer would not know that the packets could not be sent to the destination).

The routing function knows a node as **reachable** if it can access that node by a usable path. A path is usable if its path cost and path length do not exceed the limits specified in the routing node's executor database. The limits are defined chiefly by the MAXIMUM HOPS and MAXIMUM COST parameters.

In a multiple area network, a reachable area is one that the routing function can access over a path that does not exceed the AREA MAXIMUM COST and AREA MAXIMUM HOPS parameters specified in the executor database (see Section 2.5.2).

- **Distributing resources between circuits and between locally generated traffic and routing-service traffic, by:**
 - Limiting the number of packets queued for transmission on individual circuits. This prevents a circuit from becoming overloaded.
 - Managing buffers. Routing manages buffers for packets being routed to other destinations.
 - Regulating the ratio of packets to be forwarded through a node with those that are generated on that node. This achieves a balance of traffic from both sources. It prevents locally generated packets from degrading a node's routing service.
- **Tracking the number of nodes a route-through packet has visited and discarding packets that have exceeded a predefined limit** — This ensures that packets can never loop endlessly through the network tying up network resources. It helps maintain higher routing throughput by minimizing traffic and queuing delays.
- **Performing node verification** — Routing will exchange verification passwords with an adjacent node if so requested by the network management layer. This helps prevent unauthorized access from the adjacent node.
- **Maintaining counters and gathering event data for network management purposes** — This enables a network manager to identify and locate network problems such as traffic congestion.

2.4 Routing Control Messages

DECnet routing uses seven types of routing control messages to initialize the routing layer, maintain routing data, and monitor the status of adjacent nodes. **Routing update messages** are used on all types of circuits. The **Ethernet End node Hello message**, the **Ethernet Router Hello message**, and the **Designated Router Hello message** are used on Ethernet circuits only. The **Hello and Test message**, the **Initialization message**, and the **Verification message** are used on non-Ethernet circuits only.

1. **Routing Update Messages** — Adjacent routing nodes exchange routing update messages containing information about the cost and hops for each node in the network, from the perspective of the node sending the message.
2. **Ethernet Router Hello message** — This is used for initialization and periodic monitoring of routers on an Ethernet circuit. Each router on the Ethernet uses a multicast operation to periodically broadcast an Ethernet Hello message to all other nodes (routers and end nodes) on the same Ethernet. The message contains a list of all routers on the Ethernet from which the sending router has recently received Ethernet Router Hello messages. By exchanging these messages, all routers remain informed of the status of the other routers on the Ethernet.
3. **Ethernet End Node Hello message** — This message is used for the initialization and periodic monitoring of end nodes on an Ethernet circuit. Each end node periodically broadcasts an Ethernet End Node Hello message to all routers on the Ethernet. The routers use this message to keep the status of end nodes on the Ethernet.
4. **Designated Router Hello message** — This message is sent by the Ethernet designated router to all end nodes to inform them of which router should be used on the Ethernet. Designated routers and routing on Ethernet is discussed in Section 3.2.1.2.
5. **Hello and Test message** — This message tests an adjacent node to determine if the adjacency is accessible. An **adjacency** is the combination of the adjacent node and the connecting circuit. If either is not operating, then the adjacency is considered inaccessible. The routing layer of a node sends this message periodically on non-Ethernet circuits in the absence of other traffic. The Hello and Test message also checks for data integrity.
6. **Initialization message** — This message is used by the routing layer when initializing a non-Ethernet circuit. The message contains information about the node type, required verification, maximum receive-block size for the data link layer, and routing version.
7. **Verification message** — This message is sent along with the Initialization message on a non-Ethernet circuit if verification is required.

2.4.1 Segmented Routing Messages

In DECnet Phase III, the largest routing update message can handle cost and hop information for a maximum of 255 nodes. Furthermore, Phase III routers have to send complete updates containing information about all nodes, whether or not their status (reachability) has changed.

Phase IV allows segmented routing messages, which means that routing updates can be sent in multiple messages. Also, Phase IV update messages contain only the updated information. Therefore, the size of the routing messages and the number of buffers required to receive them are reduced.

2.4.2 Timing of Routing Message Transmissions

Routing messages are sent automatically when changes in the network occur. They are also sent at regular intervals to ensure that all routing nodes have up-to-date information. The time intervals are determined by timers, which are set at optimal values and should not be altered.

On non-Ethernet circuits, the timer is called the **routing timer**. When the timer expires on a node, that node sends a routing message to all adjacent nodes.

On Ethernet circuits, the timer that controls the frequency of routing updates is the **broadcast routing timer**. When this timer expires, the local node multicasts a routing message to all routers on the Ethernet. Note that if the timer is set to a value smaller than recommended, the amount of routing control traffic in an Ethernet network may become excessive.

2.5 Routing Parameters

Two basic types of parameters affect the operation of routing in Phase III and Phase IV routing nodes. **Routing control parameters** indirectly control the path that data takes through the network and also control the timing of routing messages. These parameters include:

- CIRCUIT COST
- MAXIMUM COST
- MAXIMUM HOPS
- MAXIMUM VISITS
- ROUTING TIMER
- AREA MAXIMUM COST
- AREA MAXIMUM HOPS

Routing configuration parameters specify how the network is to be configured. These parameters include:

- TYPE (Phase IV routing, Phase IV nonrouting, area router)
- MAXIMUM ROUTERS
- MAXIMUM BROADCAST ROUTERS
- MAXIMUM BROADCAST NONROUTERS
- MAXIMUM ADDRESS
- MAXIMUM AREA

This section discusses the routing control parameters that indirectly control routing. Chapter 4 discusses the routing configuration parameters.

Routing control parameters have reasonable default values for most networks. Many of the parameters can be changed after a DECnet system has been installed, such as to improve network performance or to reflect changes in the network configuration. However, any changes to these and other routing parameters should be made only after careful consideration of their effects on the local node and on the network as a whole.

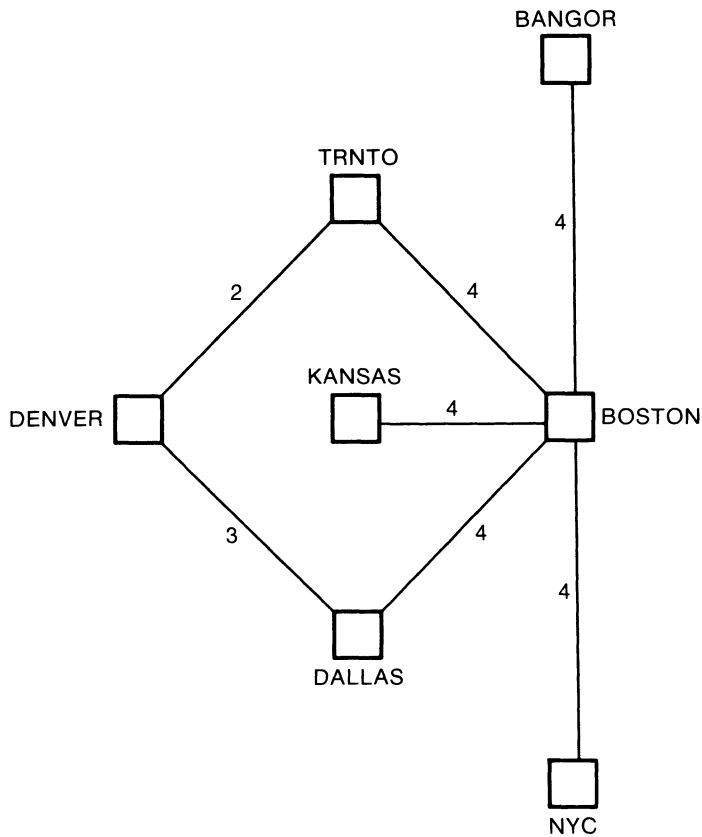
2.5.1 Circuit Cost Parameter

The COST parameter specifies the circuit cost. Because the routing layer forwards packets on the least costly path to a destination, a higher cost reduces traffic on the circuit. (Note that two nodes connected by the same circuit can have different costs assigned to that circuit.)

If a node has several lines on paths to the same destination, and one of the lines is much more expensive to use, traffic to the destination can be diverted from the expensive line by raising the routing cost of the circuit that corresponds to the line. In fact, this can be done whenever the system manager or network manager wants to decrease the traffic on a line. The manager can change circuit costs dynamically. This can be a great advantage when network performance suffers due to excessively high levels of traffic on low-cost circuits.

If the network is very large, it may be helpful to have a network manager, rather than individual system managers, be responsible for controlling the flow of data through the network.

Changing the cost of a circuit does not always change the amount of traffic it carries. For example, if a circuit leads to an end node, such as the circuit leading to end node BANGOR in Figure 2-4, the circuit cost does not affect that circuit's traffic because all traffic arriving at the end node is intended for it (an end node only has one circuit).



LKG-0560

Figure 2-4: Network Circuit Costs and Path Hops

If the line between nodes DENVER and DALLAS is expensive to use, the network manager can assign a higher cost to the circuit on that line. Then, if node DENVER has data to send to node BOSTON, it will use the lesser cost path which crosses through node TRNTO.

Circuit costs should be based on circuit bandwidth. Since packets take the path with the lowest cost, circuits with high bandwidth should normally have low costs assigned to them.

The network manager should establish a circuit cost standard that is uniform across the entire network. Section 4.1.4 gives a formula for assigning appropriate costs.

2.5.2 Maximum Path-Control Parameters

The **MAXIMUM COST** and **MAXIMUM HOPS** parameters are used to determine whether a destination is reachable. A node is reachable if there is an available path that does not exceed the value of the **MAXIMUM COST** and **MAXIMUM HOPS** parameters.

The value of the **MAXIMUM HOPS** parameter should always be equal to or greater than the longest possible path within the network. It must be less than or equal to the value of the **MAXIMUM VISITS** parameter, discussed in Section 2.5.3. The maximum cost and hops values should be chosen carefully, based on the intended use of the network, the actual network configuration, and possible failures. The default values of these parameters are reasonable for most networks.

For the network shown in Figure 2-4, a **MAXIMUM HOPS** parameter value of 6 would be sufficient. Under normal conditions, a value of 3 would be sufficient. However, a failure of the **TRNTO-BOSTON** circuit would render **TRNTO** unreachable from **NYC**, **KANSAS**, or **BANGOR**, even though a physical path still exists (the four-hop path **NYC-BOSTON-DALLAS-DENVER-TRNTO**). Consideration of possible failures is also important in determining the maximum cost.

The **MAXIMUM COST** parameter limits the possible routes a packet can take. Packets cannot take any paths whose cost exceeds the parameter's value. Thus, you may use the parameter to prevent certain paths from being used. Also, when you decrease the **MAXIMUM COST** value in the network, the network will take less time to reconfigure when a node becomes unreachable.

In multiple area networks, the **MAXIMUM COST** and **MAXIMUM HOPS** parameters are defined uniformly on an area-by-area basis. These parameters are used to limit the paths that can be taken between any pair of nodes within the same area. A second set of parameters, **AREA MAXIMUM COST** and **AREA MAXIMUM HOPS**, is used to limit the total cost and length of paths between level 2 routers within the whole network.

The **AREA MAXIMUM COST** parameter specifies the limit on the total path cost between the local level 2 router and any other level 2 router in the network. The **AREA MAXIMUM HOPS** parameter specifies the maximum number of hops that a packet can traverse between the local level 2 router and any other level 2 router in the network.

The area routing function uses these two parameters to determine whether an area is reachable. A remote area is reachable if the total path cost and path length to the nearest level 2 router in the remote area are less than the values specified for the **AREA MAXIMUM COST** and **AREA MAXIMUM HOPS** parameters, respectively. The default values for these parameters are reasonable. The selection of any other values should be made carefully, based on the level 2 (area) topology of the network.

Section 4.1.5 provides guidelines for choosing the most appropriate values for path-control parameters.

2.5.3 Route-Through Control Parameter

The MAXIMUM VISITS parameter limits the maximum number of nodes that a packet can be routed through before arriving at its destination. This is the parameter that prevents packets from being looped through the network endlessly. If the number of nodes that a packet visits exceeds the value of this parameter, the packet is discarded.

In general, use a value that is at most two times the value of the MAXIMUM HOPS parameter and at least equal to or greater than that value.

2.6 DECnet Phase IV Routing Enhancements

Two enhancements to DECnet Phase IV that affect routing are:

1. Area routing improvements
2. Path-splitting capabilities

2.6.1 Area Routing Improvements

DECnet Phase IV now differentiates between level 1 and level 2 multicast routing update messages. Previously, level 1 and level 2 routing update messages were sent to all routers; now, level 2 update messages are sent to level 2 routers only. This relieves level 1 routers from the extra processing required for the level 2 messages. The multicast addresses used for level 1 and level 2 routing update messages are:

Level 1 Routing Update Messages: AB-00-00-03-00-00

Level 2 Routing Update Messages: 09-00-2B-02-00-00

(Level 2 routers send level 1 update messages to the level 1 routing update message address.)

2.6.2 Path Splitting

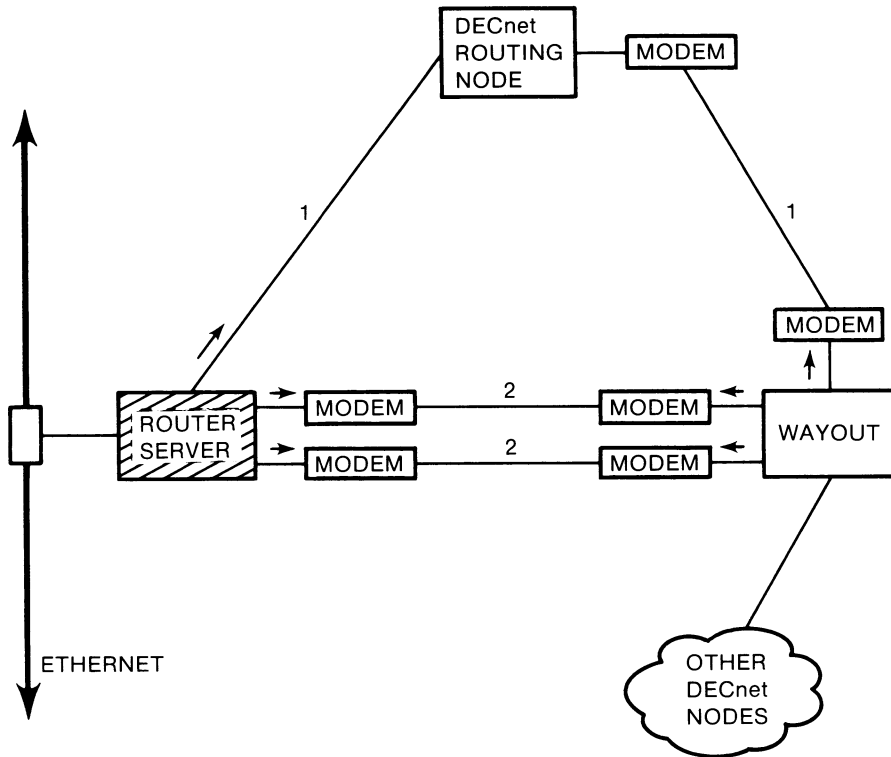
With enhancements to DECnet Phase IV, certain routing nodes can support **path splitting**. If such a routing node has multiple paths of equal cost connecting a given destination, it can divide the traffic destined for that destination evenly over the equivalent paths. This feature decreases the likelihood of any one line becoming congested. Prior to the addition of this feature to DECnet, all traffic was sent over one line only, leaving the remaining lines idle. So, besides helping prevent congestion, path splitting also provides better **utilization** of network resources.

To enable path splitting, the network manager assigns circuit costs establishing two or more paths of equal cost to a destination. The network manager also must set the PATHSPLITS parameter (on any routing nodes that will split the traffic over multiple paths) to the maximum number of paths over which the traffic will be split.

Figure 2-5 shows a simple example of path splitting. Costs are assigned to the circuits along the three paths between the router server and node WAYOUT so that the total path costs are equal (each path cost equals 2). In this way, path splitting can be enabled at the router to facilitate communications between these two systems. If traffic between the two systems is heavy, path splitting can help prevent congestion. Data will flow between the nodes without congesting any one line.

With path splitting, the destination may receive packets out of order. Therefore, the destination node, and all nodes reachable through a router that performs path splitting (such as the nodes in the remote DECnet network connected by node WAYOUT in Figure 2-5), must support **out-of-order packet caching** (also referred to as out-of-order packet reassembly).

If a destination does not support out-of-order packet caching, then the PATHSPLITS parameter can be set to 1 on routers that have multiple paths to that destination. This disables path splitting. In complex networks with nodes that do not support out-of-order packet caching, it is a good idea to set the PATHSPLITS parameter to 1 on all routers, unless you thoroughly understand all paths in the network. (Note that node WAYOUT can path-split in the opposite direction, too.)



LKG-0561

Figure 2-5: Path Splitting

Path splitting can also connect modems to remote facilities. The transmission load can be split over several low-cost, low-speed switched circuits between the router and the remote facilities, with the combined throughput equaling that attained by an expensive high-speed line.



3

Network Technologies and Topologies: Ways of Building Networks

DECnet permits great flexibility in configuring a network. A network can start with two nodes and be expanded to include over 64,000 nodes. A network manager can configure a **local area network**, with nodes located in a building or spanning a group of buildings, and expand the configuration to a **wide area network**, with nodes distributed in different cities across the country and even in other countries or continents.

A very large DECnet network (exceeding 1023 nodes) can be configured by dividing it into areas. The division of the network can be based on topology or traffic (where nodes are grouped by frequency of communication). This improves the efficiency of the network.

This chapter discusses how nodes and lines can be arranged in various configurations. It also shows how routing permits great flexibility in the way that nodes and lines are combined and connected in a network. Particular attention is given to the use of routers for building networks.

The major factors affecting network configuration are:

- Performance requirements of the applications

What are the required response time, throughput, and availability? For example, if high availability of network resources is required at many locations, then the most appropriate type of network may be a distributed network with a high degree of connectivity between nodes (multiple paths between source and destination).

A distributed network does not have a single point of failure as has a centralized network. If a node fails, the rest of the network may continue to operate with minimal impact on performance. With alternate paths available between two nodes, the obstruction of one path does not impact performance significantly.

- Characteristics of the network environment

Must the network span a limited area or a wide area? Are the locations of the nodes determined by geography or work patterns? Are the locations predetermined and fixed, or are they flexible?

- Characteristics of the nodes

Are they routing nodes or end nodes? What is the capacity of the routing nodes (how many lines are supported and what is the maximum aggregate bandwidth?) How do the nodes function?

In a network where one or a few nodes function mostly as resource providers for a relatively large number of other nodes, a centralized structure is best. Here, a central node controls the other nodes, determining which node can access the network at any time. In contrast, in a network where all nodes function both as resource providers and resource users, as in most DECnet networks, a fully distributed structure is more appropriate. In such a structure, nodes have equal "power" relative to network access. They function as peers.

As mentioned above, the functional-nature of the nodes of a network affects the performance of the network. The major performance drawback of a centralized topology or control scheme is the effect on the network's reliability. The central node is a single point of failure. Usually the central node handles most of the communications responsibilities for the network. Thus, if the load on that node becomes excessive, data transfer rates may suffer throughout the entire network.

Also, in centrally controlled networks, nodes must wait for permission from a central node before they can access the network. The wait can affect response times, especially where frequent network access is required by an application.

In contrast, distributed networks tend to be more responsive to local needs of nodes. Nodes do not have to wait for permission from a central node to access the network. With less access time, better response times can be attained over these networks. Where frequent accesses are required by an application, greater throughput can be attained.

- Characteristics of lines and circuits

What is the cost and availability of the lines? What type of circuits are supported?

These capabilities or characteristics will dictate the limits within which networks may be configured, as discussed in the following sections.

3.1 DECnet Environments

There are two basic environments for DECnet networks: (1) a wide area networking environment using traditional synchronous and asynchronous communication links and (2) a local area networking environment using high-speed, limited distance communication links. A DECnet network can include both environments simultaneously. Within these two types of environments, nodes can be configured in several ways, depending on application requirements.

3.1.1 Wide Area Networks

A wide area network (WAN) is generally used for long-distance communications. Nodes may be located many thousands of miles apart. A WAN is composed of nodes connected by individual communications links configured in various patterns. Routing plays a very important part in the configuration of WANs. It provides flexibility and reduces the need for expensive communications hardware, as explained in Section 1.2.

In a WAN, messages can be transmitted over dial-up or leased lines, or by means of **packet switched data networks (PSDNs)**. WANs typically use common carriers, like the telephone network or postal telegraphic and telephone authorities (PTTs), to transport messages over most or part of the distance. Because WANs use the telephone system for most communications, modems or modem equivalents are generally needed on each end of a communications link to convert data to analog form, and vice versa, and to permit transmission over long distances.

A network manager can specify different transmission speeds for the various communication lines in a WAN. Transmission speeds depend on the supported hardware devices and type of lines used. The hardware devices and types of lines to configure in a network depend on how much traffic they will have to handle and how much delay network users will tolerate. Specifying transmission speeds gives the network manager some flexibility in distributing workloads and adjusting performance levels for various network components.

3.1.2 Local Area Networks

Local area networks (LANs) are privately owned networks optimized for connecting information-processing equipment in a limited geographical area, such as an office, a building, or a complex of buildings. Because LANs cover limited distances, they can use materials and technologies that would be too expensive for WANs but which give very high transmission speeds and data integrity.

LANs can be designed with a variety of technologies and arranged in different configurations. Consequently, they vary significantly with respect to their transmission speeds, the distances they span, and the capabilities and services they offer.

An Ethernet **baseband** LAN consists of a segment of coaxial cable or several segments of coaxial cable joined together, with each segment ranging in length from 20.4 to 500 meters. The Ethernet cable is terminated at both ends. The cable can support bit rates as high as 10 Mbps.

An Ethernet **broadband** LAN uses a different cabling system than baseband LANs. The cable can be divided into many frequency ranges such that different varieties of video, voice, and data equipment can transmit over the channel simultaneously.

A single-area DECnet Phase IV Ethernet network supports up to 1023 nodes. It supports a **bus topology**, a single communications medium to which all the nodes are attached as equals. The single network cable replaces the numerous interconnecting cables usually required in traditional WANs.

Segments of coaxial cable can be connected to extend an Ethernet LAN beyond the 500-meter limit of a single segment. Extended LANs not only introduce larger networks in terms of distance but also in terms of the number of nodes supported. Cable segments are joined using a device called a **repeater**. The repeater enables the connected segments to function as if they were one cable. Ethernet LANs can be further extended by **bridges**, which are more sophisticated than repeaters. Bridges are discussed in more detail in Section 3.5.

All DECnet nodes directly connected to an Ethernet cable must be Phase IV nodes. However, Phase III nodes or Phase IV nodes that are not on the Ethernet can gain access to Ethernet resources through a router on the Ethernet. The router can be a server-based router or a full-function Ethernet node acting as a Phase IV router. Routers can be used on an Ethernet to provide connections between the Ethernet nodes and WANs or other LANs.

Routers are not required on an Ethernet for communications among nodes on the same Ethernet. Every node has a direct link with every other node on the same Ethernet. All nodes on the Ethernet are considered adjacent, regardless of whether they are physically adjacent. Routers are required on an Ethernet for communications with nodes off the Ethernet. They may also be required when the Ethernet is divided into areas, for communications with nodes in other areas on the Ethernet. (If no routers are present, then all nodes on an Ethernet LAN can communicate with each other, even across areas. However, if there are any level 1 routers on the LAN, then level 2 routers are needed for communications across areas.)

Ethernet nodes can communicate with packet-switched data networks (PSDNs) that implement the X.25 protocol and with systems in IBM SNA networks. To permit Ethernet nodes access to these networks (and vice versa), Digital Equipment Corporation offers **gateways**. The gateways can be software packages that reside on one of the nodes on the Ethernet, or they can be server-based units that connect directly to the Ethernet.

3.2 Data Link Technologies

DECnet supports several types of network connections:

- A connection to an Ethernet circuit in a LAN configuration
- A connection to a node running DECnet using the Digital Data Communications Message Protocol (DDCMP), being either a synchronous point-to-point or multipoint connection, or an asynchronous point-to-point connection
- A connection over the computer interconnect (CI) to another node running DECnet
- An X.25 connection to a node running DECnet

This section focuses on the first two types of connections: Ethernet connections and DDCMP connections.

Network nodes in a WAN can be connected using DDCMP point-to-point or multipoint lines and circuits. Ethernet LAN nodes are connected by means of the bus topology.

A **point-to-point** line connects two nodes using a single circuit. A **multipoint** line is shared by more than two nodes. Each node communicates over the line using a separate circuit. One node on the line, called the **control station**, controls access to the shared communications path. The other nodes are called **tributaries**.

The type of line or circuit between two nodes can have a marked effect on access capabilities. For example, nodes connected by a point-to-point line/circuit always have access to the communications path linking them together. In contrast, the tributaries on a multipoint line cannot access the path until **polled** by the control station (tributaries are usually polled on a round-robin basis). Thus, a tributary has to wait its turn before being able to access the communications path. Also, tributaries always have to send data to the control station, which forwards the data to other tributaries, if necessary. Point-to-point nodes can communicate directly with other nodes.

Data is usually transmitted at similar speeds over point-to-point and multipoint connections, but point-to-point connections provide higher throughput. The difference is that each station on a multipoint line has only a certain proportion of the available bandwidth, and the control overhead is very high.

Nodes on an Ethernet can access the network any time the Ethernet line is clear (that is, when no other node is transmitting at the same instant, as explained in Section 3.2.1.1). Because data is transmitted over the Ethernet at such high speeds, transmissions complete in a relatively short time. This leaves the line virtually clear all the time except when traffic demands are unusually high.

Transmissions over an Ethernet are faster than over point-to-point connections.

3.2.1 Ethernet Connections

At the data link layer, network control for an Ethernet is multiaccess and is fairly distributed to all nodes. In other words, the Ethernet is a single, shared channel with many nodes that have equal rights to gain access to it. The process that fairly distributes and regulates access to the channel is called **Carrier Sense, Multiple Access with Collision Detect (CSMA/CD)**.

3.2.1.1 Access Control — With CSMA/CD, each node listens to the channel and holds transmission until no signals are detected (carrier sense). Once a node proceeds to transmit, it can detect if another node has also begun transmitting at the same time (collision detect). When the transmissions of two or more nodes overlap (collide), the transmitted messages become unusable and must be retransmitted. So, any node that detects a collision will stop transmitting and wait for a randomly selected, short time before retransmitting.

Each node on an Ethernet can detect every message transmitted by other nodes. Some messages are intended for all nodes (**broadcast address**), some are intended for a subset (**multicast address**), and some are intended for individual nodes (**physical address**). Because every node can detect every message, and messages can be addressed to their intended recipient(s), greater communications efficiency is attained than on a completely connected DDCMP network.

3.2.1.2 How Routing Works on an Ethernet LAN — All nodes on an Ethernet can communicate directly with each other without depending on intervening routing nodes. Initially, however, end nodes on an Ethernet do not have information about other nodes on the network. When there is no router on the Ethernet circuit, the end node trying to communicate with another node transmits the message on the Ethernet channel, assuming that the other node is there. The message travels along the Ethernet channel until it reaches the destination node (the destination node can recognize and retrieve messages addressed to it).

If the destination node *is* on the Ethernet, it informs the source node. The source node then adds the destination node to its **cache** and sends subsequent messages directly to the destination node as if there were a point-to-point connection between the two nodes.

If the destination node is *not* on the Ethernet, communication will not be possible. End nodes on the Ethernet must use a router to communicate with nodes off the Ethernet.

If one or more routers are on the Ethernet, end nodes will initially know one of the routers as the **designated router**. This is a node designated to provide message routing services for end nodes on the Ethernet. End nodes can use the designated router when communicating with any node for the first time.

If the source node, destination node, and designated router are all on the same Ethernet, the first message between the source and destination will be sent to the designated router. When the designated router detects that both the source and destination

nodes are on the same Ethernet, it will inform the destination node while routing that first message to it. The destination node will then send any traffic for the source directly to the source, bypassing the router. Thereafter, the source and destination nodes will communicate directly.

The process is similar for an Ethernet node trying to communicate for the first time with a non-Ethernet node, reachable through an Ethernet routing node other than the designated router. The designated router will cause all subsequent traffic to flow directly between the Ethernet node and that Ethernet routing node.

If two or more routers are on the same Ethernet, the network software on all routing nodes on the Ethernet elects one of them as the designated router. The election is made on the basis of the highest numerical priority, a circuit parameter called the ROUTING PRIORITY parameter. The system or network manager can define this parameter for each Ethernet circuit on each routing node.

If the values for the highest priority router are the same, the router with the highest node address is elected as the designated router.

3.2.2 DDCMP Connections

DDCMP is a protocol designed to provide an error-free communications path between adjacent nodes. DDCMP is used for multipoint and point-to-point connections. Today it is used most often for point-to-point connections.

Point-to-point circuits and multipoint circuits perform as virtual circuits: nodes on these circuits interact as though a specific circuit were dedicated to them throughout the transmission (on multipoint circuits, however, the actual physical connection is allocated by DDCMP on the control station). Also, individual nodes on DDCMP circuits must be addressed directly; no multicast or broadcast addressing capability is available as with Ethernet.

Point-to-point connections are either **synchronous** or **asynchronous**. Synchronous transmissions are more efficient when large blocks of data are being sent. Bandwidth is saved and higher throughput can be achieved. Synchronous signaling is predominant at line speeds above 19.2 kbps and over lines provided by PTTs. It is usually used over long-haul connections, such as crossing a country or ocean. (The Ethernet is used for high-speed connections over short distances.) Asynchronous signaling is usually employed on lower-speed or private (cheaper) lines. For example, it provides an inexpensive way of connecting a MicroVAX system to a VAX 8600 system.


The DDCMP protocol can be used for local synchronous or asynchronous connections as well as for remote synchronous or asynchronous connections over a telephone line using a modem.

Digital Equipment Corporation offers Ethernet routers that provide synchronous and asynchronous connections to nodes off the Ethernet. A synchronous router is used for high-speed connections, such as to VAX 8600s or to another router that connects to another LAN. An asynchronous router provides relatively inexpensive low-speed connections to smaller, less-demanding devices, such as personal computers and other systems for which high-cost connections would be unsuitable.

Figure 3-1 shows several synchronous server-based routers being used on an Ethernet LAN. The routers RTRSV1 and RTRSV2 are used to connect LAN A and LAN B. A single, high-speed synchronous communications line between the two servers will handle a large amount of traffic. Without routers RTRSV1 and RTRSV2, LAN B was a distinct network disconnected from the network formed by LAN A and the WAN. By adding these two routers, full connectivity is attained. Both LANs and the WAN become one extended network.

Though bridges are usually used to connect two LANs, routers may be used in the configuration shown in Figure 3-1 because the synchronous line can span longer distances. Section 3.5 discusses other reasons for using routers instead of bridges, and vice versa.

Note that in Figure 3-1, nodes on LAN B also have access to the WAN by means of the third router, RTRSV3. RTRSV3 has three synchronous lines connecting two large systems and the WAN.

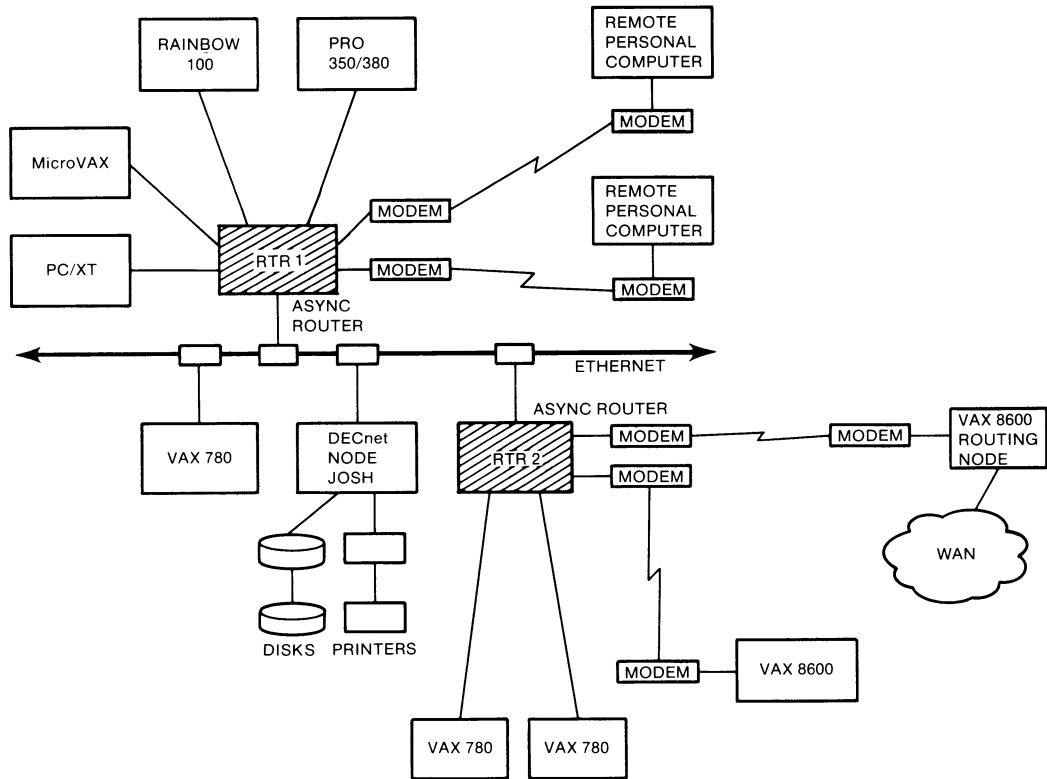


LKG-0562

Figure 3-1: Synchronous Routers

Figure 3-2 shows two asynchronous server-based routers: nodes RTR1 and RTR2. RTR1 connects several Digital personal computers plus an IBM PC/XT to the Ethernet. It also connects two remote personal computers via telephone dial-up lines. RTR2 connects several larger systems and a WAN to the Ethernet. Asynchronous connections are used here perhaps because the communications traffic is minimal, not requiring higher-cost synchronous connections.

Any node directly or indirectly connected to either router in Figure 3-2 has full access to the entire network shown. For instance, all the resources of node JOSH are accessible regardless of the type or location of the connection to the router servers.



LKG-0563

Figure 3-2: Asynchronous Routers

3.3 Multiple Area Networks

As network manager, you can set up a large network of more than 1023 nodes by dividing it into areas. Area routing permits configuration of networks of up to 64,000 nodes.

You can take an existing large network and partition it into areas, or you can take existing networks, designate them as areas, and combine them to form the larger multiple area network.

Besides letting you configure very large networks in an efficient manner, multiple area networks make it easier to merge existing networks. The nodes do not have to be renumbered. An area number is added to the number of each node to keep their identification unique.

Area routing can also be advantageous for certain networks of less than 1023 nodes. For example, if you have a network of 1000 nodes, the routing overhead incurred by each routing node can lead to inefficiency. By dividing your network into areas, such as five areas of 200 nodes each, the overhead for routing between areas will be restricted to level 2 routers only. With a reduction in routing traffic, your network will gain in efficiency. Also, by dividing your network into areas, it will be easier to manage.

As another example of how area routing can increase the efficiency of a network of less than 1023 nodes, suppose you have such a network and many of the nodes are DECnet Phase III. Because DECnet Phase III nodes have an addressing limit of 255, many nodes will be unable to communicate with each other. You can divide the larger network into groups (areas) of nodes, each group containing fewer than 255 nodes. Keeping in mind that the Phase III nodes can communicate with nodes in the same area only, assign to each area those nodes that need to communicate with each other.

To ensure that all the Phase III nodes within an area can communicate with each other, renumber all nodes in the area so that all node numbers are less than 255. Section 4.3 provides guidelines you should follow when setting up a multiple area network. The guidelines also provide suggestions for configuring Phase III nodes within a multiple area network.

To add area support to a standard network, follow the steps outlined in Section A.1 of Appendix A. Section A.2 discusses the steps to take when designing a multiple area network from the beginning. Where possible, avoid setting up Ethernet LANs as multiple area networks, especially in large networks. Section 5.2 discusses problems to avoid when configuring multiple area networks.

3.4 Node Characteristics

Characteristics of a node that affect network configuration are whether the node supports DECnet Phase III or Phase IV, whether it is a routing node or end node, and whether it is a dedicated or general purpose node.

3.4.1 DECnet Phase III and Phase IV Nodes

In planning a network configuration, the phase of DECnet that is installed on a node is important. DECnet Phase IV extends the networking capabilities of a node beyond what is offered in Phase III. Only Phase IV nodes can be connected directly to an Ethernet cable. Phase III nodes can communicate with Ethernet Phase IV nodes through a routing node on the Ethernet.

DECnet Phase IV is backward-compatible with Phase III, meaning that DECnet Phase IV networks can coexist with Phase III networks. Utilities and services that operate among Phase III products can operate between Phase III and Phase IV products. Most programs developed for Phase III will run on Phase IV. Phase III nodes can communicate with Phase IV nodes directly, or by way of a router.

Phase IV supports routing in Ethernet LANs of 1,023 nodes and area routing in networks of over 64,000 nodes, while Phase III supports up to 255 nodes. Therefore, in configuring a network of both Phase IV and Phase III nodes, consideration must be given to how these different types of nodes are interlinked. Sections 4.1.2 and 4.3 provide guidelines. Section 5.2 discusses problems that can develop if Phase III and Phase IV nodes within the same area of a network are not configured properly.

Communications servers (such as server-based routers) all implement Phase IV software. These servers are important to configuration plans because they off-load functions, such as routing, that are normally performed by a full-function node.

3.4.2 Routing Capabilities

There are two types of DECnet nodes: end nodes and routing nodes. End nodes cannot receive and forward messages intended for other nodes. They can only send messages to an adjacent node. If an end node has multiple circuits to one or several adjacent nodes, only one of those circuits can be active at a time.

If a node adjacent to the end node is a routing node, its routing capability can be used by the end node to communicate with other nodes in the network.

A routing node can send or forward messages to any other Phase IV nodes in the network, regardless of their location. It can receive packets from any other Phase IV nodes. It can have multiple circuits actively communicating with one or several nodes simultaneously. Routing nodes maintain an up-to-date routing database and determine the best paths for sending or forwarding data to a destination node.

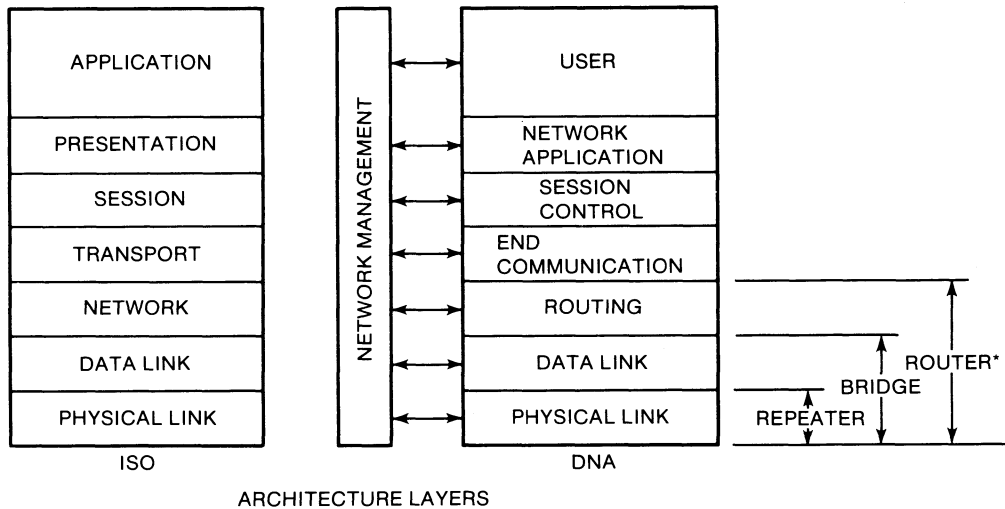
3.5 Routers and Other Products That Extend Ethernet LANs

This section discusses several products used to extend Ethernet LANs, focusing particularly on routers and bridges. Table 3-1 compares the characteristics of routers and bridges. Figure 3-3 shows how repeaters, bridges, and routers differ with respect to the layers of the network that they act upon. The layers of both the Digital Network Architecture (DNA) and the International Standards Organization (ISO) architecture are shown.

Table 3-1: Comparison of Routers and Bridges

Characteristic	Routers	Bridges
Network Environment	WANs or LANs	LANs
Multi-protocol	No	Yes
DNA Layer	Routing	Data Link
Terminal Communications		
LAT support	No	Yes
CTERM support	Yes	Yes
Network Management	High	Medium
Path Splitting	Yes	No
Area Routing	Yes	No
Reliable Data Service	Yes	No
Message Delay	High *	Low
Performance (Throughput)	Low *	High
Cost	Medium	Low to Medium

* Keep in mind that the ratings of router message delay and performance are "inferior" only in relation to bridges. Under normal conditions, the delay incurred by routers is minimal and imperceptible by network users and the throughput exceeds user requirements.



*THE ROUTER CAN FUNCTION AS A NODE IN THE NETWORK, ADDRESSABLE BY OTHER NODES FOR NETWORK MANAGEMENT FUNCTIONS AND OTHER PURPOSES. THEREFORE, ITS FUNCTIONAL SCOPE ACTUALLY SPANS THE LAYERS ABOVE ROUTING.

LKG-0564

Figure 3-3: Functional Scopes of Repeaters, Bridges, and Routers

3.5.1 Bridges

A standard Ethernet LAN has a distance limitation of 2,800 meters (9,194 feet) between the furthest two nodes on the network. This limitation is the sum of several components including two 50-meter (164 feet) transceiver cables, three 500-meter (1,640 feet) coaxial cable segments, four repeaters (including four 50-meter transceiver cables), and 1,000 meters (3,280 feet) of fiber optic cable used for expanding the network (with remote repeaters).

A maximum of two repeaters can be placed in the path between any two transceivers. A **repeater** is a simple hardware device used to connect one Ethernet segment to another. A repeater retimes, amplifies, and relays all signals it receives from one segment to the next. A repeater does not perform any of the data link functions such as addressing, error control, or flow control.

A bridge can be used to extend a LAN beyond the limitations mentioned above. An extended LAN may consist of up to 8,000 nodes and, using fiber optic bridges, can span distances of up to 22 kilometers (13.5 miles). Up to seven bridges can be used to connect eight LANs.

A bridge serves as a relay on the data link level, the second layer of the DNA model. A bridge performs several important data link functions. To prevent two LANs from being overwhelmed by their combined traffic, the bridge controls network traffic between the LANs it connects. It selectively forwards packets to keep local traffic local. (Local traffic is traffic transmitted from the source LAN whose destination is the source LAN.) Only information destined for different LAN segments passes through the bridge and continues on to the appropriate remote destination.

This ability to filter packets reduces traffic over the entire network, increasing the network's bandwidth efficiency. To overcome bandwidth limitations of a LAN, a bridge automatically isolates high-volume areas from the rest of the network.

Because bridges operate at the data link level, they are **protocol transparent**, meaning that they allow various higher-level protocols to coexist on the same wire. Therefore, along with DECnet higher-level protocols, bridges will allow TCP/IP, XNS, Local Area Transport (LAT) and other protocols. (However, they do not allow incompatible protocols to intercommunicate. To translate between different higher-level protocols, such as SNA and DECnet, gateways are required.)

3.5.2 Routers

Routers operate on the routing level, which is "above" the data-link level on which bridges operate. Thus, routers provide more services than bridges, but they are limited to DECnet protocols. They do not transmit packets that vary from the DECnet protocols.

While bridges are (ideally) transparent to the nodes communicating through them and appear as if they were just part of the Ethernet wiring, routers are visible as individual nodes in the network. A router can be directly addressed by nodes using it to forward packets.

Routers differ from bridges in several other ways. Routers are used to connect LANs to WANs, while bridges are used in LANs only.

As network supervision devices used in both LANs and WANs, routers allow nodes to locate one another. Also, routers are needed in each area of a multiple area network.

Besides performing routing and all the tasks related to that service, including path splitting and area routing, routers can optimize the path from a source to a destination. They direct packets only along those paths that lead to their destinations, and they do not forward any packets having "unknown" destination addresses (that is, node addresses that are undefined in the router's node database). Routers cooperate to manage the flow of traffic and to avoid loss of traffic. Bridges do not perform these services.

Another way in which routers differ from bridges is that routers monitor the status of the DECnet network, automatically generating alarm and status messages (event messages) when changes occur in the network. Also, routers are more tolerant of high error rates on communications lines than are bridges, offering higher data transmission reliability. Routers have greater network management capabilities than do bridges.

These differences are major factors for determining which product is more suitable for a particular application.

3.5.3 When to Use Bridges and Router Servers

Whether to use a bridge or router server to extend a LAN depends on the applications environment. Fortunately, there are guidelines that help determine which environments are more suitable for bridges or for routers. These guidelines focus on three primary considerations:

- 1. Protocol Considerations: Is the network primarily DECnet or is it multi-protocol?**

If all the network protocols are DECnet or primarily DECnet, then routers are preferable to bridges, providing that the price/performance ratios are comparable. Routers have the advantage here because they provide routing functions that bridges do not.

In a multi-protocol environment, bridges may offer a better solution than the router because they provide protocol transparency. However, if DECnet is the only protocol in the multi-protocol environment that needs extension past the local LAN, then routers are better than bridges. Routers will direct the traffic to the appropriate destinations.

- 2. Bandwidth Efficiency Considerations: Is the network small or large?**

If the size of the network is small or medium, bridges use the Ethernet bandwidth more efficiently because of their ability to isolate (filter) traffic.

In larger, more complex networks, routers use the communications bandwidth more efficiently than do bridges. The main advantages provided by routers in such networks is their ability to perform path splitting and to optimize the path between a source and destination. Routers direct packets to their destinations over the best paths. Bridges do not differentiate paths.

Bridges are more effective in simple network topologies. In more complex topologies, routers are necessary because of their area-routing capability.

3. **Data Transmission Reliability: Do the communications media have high error rates? Does the networking environment require increased network control?**

In large, complex networks with many levels of subnetworks, or in WANs that use leased lines with high data-error rates, routers provide the additional features required to ensure a high level of network performance and reliability. Routers know about the network topology and can detect when a node is not reachable. Thus, they will not forward data to unreachable nodes. In contrast, bridges continue forwarding data regardless of the status of destinations.

In large extended LANs, however, bridges place less burden on the routing algorithm, since all the nodes appear to bridges as end nodes. With too many routers in an extended LAN, an excess of routing overhead can hurt network performance. (All the routers are busy exchanging routing messages.) The appropriate mix of routers and end nodes in an Ethernet LAN is necessary to ensure good network performance (see Section 4.2).

4. **Distance Considerations: What are the distances to be spanned?**

Bridges are best to use for spanning short distances of less than 1000 meters (3,280 feet). Routers are best for distances greater than approximately 4800 kilometers (about 3000 miles) or for satellite connections.

The above discussion helps determine where to use either a router or a bridge. Keep in mind that for connecting two Ethernet LANs, the bridge will provide higher performance. But note that to connect three or more LANs, routers must be used.

The main difference between bridges and routers is that bridges are strictly store-and-forward devices used for connecting Ethernet LANs, while routers perform more services, connecting an Ethernet LAN to nodes off the LAN or interconnecting nodes within a point-to-point WAN. Throughput is higher for a bridge because it is only operating on the data link level. A router uses the routing layer and routing algorithms and so consumes more communications bandwidth and memory. However, routers direct traffic to the proper destinations over the best paths.

Often situations present themselves where routers and bridges are used in combination. This is particularly true in large networks with a number of areas. Routers can be used to off-load the routing from other nodes, while bridges can be strategically placed to improve network performance. Do not use a single router in parallel with a bridge (or repeater) to connect two Ethernet LANs. The bridge can handle all the traffic of both LANs sufficiently, so there is no need to have the router in parallel with it. Having a router in parallel with a bridge can cause complications affecting the router and, eventually, the whole network.

NOTE

When configuring routers in a network with bridges (or repeaters), precautions must be taken with regard to the values assigned to parameters on the router. See Chapter 4.

Strategically placed bridges in a large network can improve network performance over an all-router network because of the reasons cited above.

4

Configuring Network Nodes for Optimal Routing Performance

This chapter discusses guidelines that system and network managers should follow in configuring nodes in a network. These guidelines will help bring about the smooth operation of routing and other network functions.

Section 4.1 gives general guidelines that apply to both wide area and local area networks. Section 4.2 gives guidelines that apply to Ethernet LANs. Section 4.3 gives guidelines that apply to multiple area networks.

To configure a DECnet network, system managers must set up their systems to support the network. In other words, the system manager configures the network from the perspective of the local node's network operation. This involves supplying information to the local node defining characteristics of the node and determining how the node functions in the network (for example, whether it is a routing node or an end node). The information also defines the characteristics of various network components such as circuits, lines, and the remote nodes with which the local node will communicate. This information constitutes the **network configuration database** for the local node. Each node in the network has such a database.

A system manager should make sure that the characteristics or parameters defined in the database correspond to the needs and characteristics of the network. Also, values of certain parameters must be consistent throughout the network.

To ensure a smooth-running network, all system managers should exchange information regarding their nodes that is appropriate to network operation. The network manager may act as a focal point to ensure consistency (where required) and suitability of configuration databases throughout the network. For example, the network manager can ensure that all node addresses are unique and that routing control parameters provide efficient data flow through the network.

To assist the network manager in achieving and maintaining good network performance, system managers of network nodes should closely monitor the network-related aspects of their specific nodes.

System managers are responsible to both local and remote users. They should be aware that while local applications usually demand the greatest share of system resources, the remote users can be a very large group and must be assured of each node's response to network applications. For example, the system manager must make sure enough network buffers are available to support network communications.


Digital Equipment Corporation provides several software tools to monitor the network, such as:

- **Event logging:** The event-logging facility can be used on each node to continuously monitor the network. Whenever significant events occur in the network, they are recorded at a specified location where you can examine them.
- **Network Control Program (NCP):** This program provides commands you can use to display information about network components, test network components, and change configuration parameters. These commands can be used to solve most communications problems.
- **NMCC/DECnet Monitor:** The Network Management Control Center/DECnet Monitor is a set of sophisticated tools for the observation and control of complex networks. You can command the tools to present color graphics displays that show the condition of the network. The tools also work with a database of configuration information and network parameters; you can access and analyze the information and tune the network. The NMCC/DECnet Monitor is a layered product that runs on most VAX/VMS systems.
- **VAX ETHERnim (Ethernet Network Integrity Monitor):** VAX ETHERnim is a network maintenance application program for Ethernet. It tests the communications path to nodes on the Ethernet. It also maintains a database containing information about each node.

4.1 General Guidelines for Configuring Networks

Follow these general guidelines when setting up nodes in LANs or WANs.

1. All nodes in the network must have a unique DECnet node address and name (see Section 4.1.1).
2. All routers within an area should have the same value specified for the MAXIMUM ADDRESS parameter (see Section 4.1.2).
3. All nodes should use the same network buffer size (see Section 4.1.3).
4. The standard used for assigning circuit costs should be based on bandwidth and should be uniform throughout the network (see Section 4.1.4).


- 
5. Path lengths between nodes should be kept as short as possible, especially on heavy traffic routes (see Section 4.1.5).
 6. Traffic should be distributed, where possible, so that excessive traffic does not converge on single points in the network (see Section 4.1.6).

4.1.1 Specifying Node Identification

Each DECnet node must have a unique node address and a unique node name. Node addresses are used by DECnet software to route messages. They are known network-wide by the routers. If node addresses are not unique, communications problems may result (Section 5.1.1 discusses one such problem). To prevent duplication, node addresses should be assigned from a central registration point managed by the network manager.

On each node, define the node names and addresses of remote nodes with which the local node will communicate. Node names are known only to the local network software of each node.

To avoid potential confusion, be sure that the node names and addresses are consistent throughout the network (in other words, make sure each node is known to every other node by the same name and address). As your network increases in size, network maintenance is much easier if all nodes reference remote nodes by the same name. Again, it may be best if a central registration point assigns node names for your entire network.



In networks with multiple areas, each area number must be unique within the network and each node number unique within the area.

Because the size of a router's routing database depends on the highest node number in its area, you should assign node numbers sequentially, beginning with 1. This helps minimize the size of the routing database on routing nodes. You should also assign area numbers sequentially.

4.1.2 Assigning the Maximum Node Address on Nodes

All routers within an area of a multiple area network, or within a single area network, should have the same value specified for the MAXIMUM ADDRESS parameter (also called maximum node number).

The value used for the MAXIMUM ADDRESS parameter should be greater than or equal to the largest node number within the area. If the maximum value defined on a routing node is too small, nodes with numbers exceeding the value will be excluded from the node's routing database, and will therefore be unreachable.

You should choose a value for this parameter sufficiently larger than the current network size to leave room for future expansion. (This will depend on the anticipated growth rate of your network.) This saves you from having to change this parameter every time you add a new node to the network.

Do not set the MAXIMUM ADDRESS parameter too high. This will waste memory in nondedicated routers (routing messages will become very large) and increase the amount of routing traffic in the network.

This consideration also applies to the MAXIMUM AREA parameter defined on level 2 routers in multiple area networks. This parameter limits the number of areas that a level 2 routing node will recognize. All level 2 routers in the network should have the same maximum.

When Phase III nodes are included in a network with Phase IV nodes, it is recommended that the MAXIMUM ADDRESS parameter be given the value of 255 on all nodes. (In a multiple area network, this applies to all nodes in the same area as the Phase III nodes.) However, the MAXIMUM ADDRESS parameter can still be greater than 255 on Phase IV nodes. Nodes with an address exceeding 255 will not be reachable through Phase III routing nodes.

If the MAXIMUM ADDRESS parameter is less than 255 in a network consisting of both Phase III and Phase IV nodes, and some addresses exceed 255, then routing nodes may not receive routing update information about those nodes having addresses exceeding 255. (You will see "Partial routing update loss" event-logging messages.)

On Phase III nodes, the largest acceptable value of the MAXIMUM ADDRESS parameter (less than 256) can be computed by using the following formula:

$$\text{MAXIMUM ADDRESS} = \frac{(\text{Network Buffer Size} - 9)}{2}$$

4.1.3 Selecting the Buffer Size and Segment Buffer Size on Nodes

On all routers, the network buffer size, which is defined by the BUFFER SIZE parameter, must be equal to or greater than the value of the SEGMENT BUFFER SIZE parameter defined throughout the network. This will ensure that the buffer size will be able to accommodate the largest messages to be routed through the network.

If a router's network buffer size is too small to accommodate any incoming data packets to be forwarded across the network, then those packets will be discarded. In this case, applications that send small packets (such as remote terminal applications or most network management commands) will work fine, but applications that send large packets (such as file transfers) will "hang" (a transfer will take an indefinitely long time to complete, if it completes at all).

On a routing node, the network buffer size should be equal to the largest buffer size in the network. For best results, buffer sizes and segment buffer sizes of all nodes in the network should be equal. However, end nodes can have a smaller buffer size if the communications link to a router cannot reliably convey a large message (such as if the line is error prone). High-level protocols will then ensure that shorter packets are sent over such links.

If you are changing the buffer size throughout the network, make sure the segment buffer remains less than or equal to the buffer size. To increase the network buffer size, first increase the buffer size throughout the network, then the segment buffer size. To decrease the network buffer size, first decrease the segment buffer size throughout the network, then the buffer size. For more information, refer to the appropriate DECnet system documentation for the node.

Buffer sizes should be based on several factors:

- The record-blocking factors (transmission segment size) used by application programs.

If possible, select a buffer size (adding to it the DECnet header size of 64 bytes) that is an exact multiple of the record size used by application programs. A good size is 576 bytes. This size accommodates one disk block plus the required DECnet overhead.

- The error rate of the communications lines.

If any lines are prone to errors, specify a smaller buffer size (256 bytes, for example) to reduce the probability and cost of retransmissions over those lines.

- Your throughput requirements.

High-speed lines perform better with large buffers and large message segments because fewer messages need to be processed. This assumes the applications require the transmission of sufficient blocks of data to fill larger buffers. This may not apply to lines that are error prone, as mentioned above.

A large buffer size is also good for bulk data transfers. Consider this if a routing node will be processing a high volume of network file transfers. However, note that a large buffer size reduces the number of buffers available on some systems. On such systems, processes may have to wait for buffers to become available.

4.1.4 Assigning Circuit Costs

The network manager should establish a circuit cost standard that is uniform across the entire network. Use the following formula to determine appropriate circuit costs. The formula is based on circuit bandwidth. The bandwidth of a circuit determines the circuit delay: the greater the bandwidth, the smaller the delay, and vice versa. The formula assigns higher costs to those circuits that contribute the greatest delays.

Cost =

- 1 If the bandwidth is greater than 100 kbps
- x Where x is approximately equal to 100,000 divided by the bandwidth, and where the bandwidth is greater than 4 kbps but less than 100 kbps
- 25 Where the bandwidth is less than 4 kbps or the circuit is an X.25 circuit over a public network

4.1.5 Choosing Path Lengths

Where possible, path lengths (in terms of hops) should be kept to a minimum, especially for those paths that will be subjected to heavy traffic. The maximum path length between any two nodes depends on the response time requirements of the network users and applications. For example, for most file-transfer applications, a maximum of six hops may be reasonable. For remote terminal applications, this maximum may be too large.



By minimizing path lengths in a network, you gain better control of network traffic flow. You can set lower circuit costs, since the path costs will be less when paths are shorter. In Ethernet LANs, high path costs can cause the network to take much longer to reconfigure its routing databases when topology changes occur (such as a node failing).

However, longer paths may be necessary for economy. On a longer path, expensive communications lines and other resources can be shared and, therefore, they get more use per cost. Multiple-hop paths require less hardware expense for creating a path and afford greater flexibility than do low-hop paths. However, multiple-hop paths incur greater delays in communication than do low-hop paths. This is because of the accumulated transit delays within nodes and on any low-speed communications lines contained within the multiple-hop path.

Figure 4-1 illustrates some of the advantages and disadvantages of each type of path. With the longer path shown, messages sent from node A to node D will incur longer delays than those sent from A to D on the low-hop configuration.

In the low-hop configuration, node B incurs a greater load because of greater routing responsibilities. It has to route messages between nodes C and D, as well as between node A and nodes C and D. The low-hop configuration offers the advantage of configuring more nodes as end nodes (nodes A, C, and D). In the other configuration, only nodes A and D are end nodes.

The ideal configuration is that afforded by Ethernet LANs where all nodes can be configured as end nodes and each node is only a single hop away from every other node on the same LAN.



LKG-0565

Figure 4-1: Low-Hop vs Multiple-Hop Paths

4.1.5.1 Defining the MAXIMUM HOPS and MAXIMUM COST Parameters — With the MAXIMUM HOPS and MAXIMUM COST parameters, DECnet allows you to control the length of paths used by a router. Define these parameters as appropriate for your network configuration, keeping in mind the following considerations.

If the MAXIMUM HOPS parameter is set too low on a router, then packets that travel long distances through the network will be discarded before they arrive at their destination. This problem will cause some nodes to be unreachable if they are physically distant.

If the parameter is set too high, certain topology changes may cause an excess of traffic in the network. For example, when a node fails, packets that were in transit for that node will loop through the network until the value of the MAXIMUM VISITS parameter is exceeded. The value of the MAXIMUM VISITS parameter depends on the value of the MAXIMUM HOPS parameter. So, if the MAXIMUM HOPS parameter is set to a high value, the maximum number of visits that packets are allowed will also be high. The packets that are looping through the network will be able to continue looping for longer times before being discarded. This will consume more resources on routing nodes.

If the MAXIMUM COST parameter is set too low on a router, then packets that travel long distances (over paths with high path costs) through the network will be discarded before they arrive at their destination. This problem will cause some nodes to be reachable only from nodes which are physically close to the local router.

If the parameter's value is set too high, the network will be exchanging more routing information and will take longer to reconfigure its routing databases when a topology change occurs.

4.1.5.2 Defining the AREA MAXIMUM COST and AREA MAXIMUM HOPS Parameters — The AREA MAXIMUM COST and AREA MAXIMUM HOPS parameters must be defined on all level 2 routers in a multiple area network.

A level 2 routing node uses the AREA MAXIMUM COST parameter to control the total path cost between itself and any other level 2 router. If the path cost to a level 2 router exceeds the value of this parameter, that level 2 router is unreachable.

Select a value large enough to include all the areas that you want a routing node to be able to reach. If you have an extremely large network that includes areas with which communications will be infrequent and unnecessary, use this parameter to exclude these areas from the routing node's reachable network.

The AREA MAXIMUM HOPS parameter sets the maximum number of hops that a message can make between the local level 2 router and any other level 2 router. If the path length to a level 2 router exceeds the value of this parameter, that level 2 router is unreachable.

The value of this parameter should be large enough to include all the areas that you want the level 2 router to be able to reach.

4.1.6 Distributing Network Traffic

Where possible, network traffic should be distributed evenly throughout the network, or at least so that excessive traffic does not converge at particular points of the network (this pertains especially to routing nodes, since they are points of convergence in a network).

To prevent too much data from converging at any one node in the network, the topology can be designed so that traffic between opposite ends of the network is dispersed over several routes rather than over one or two routes with many hops. The use of path splitting can help distribute traffic loads (see Section 2.6.2).

Distributed processing is another help in this regard. It helps reduce the likelihood of congestion. With centralized processing, traffic is more likely to converge on the central processing point.

If a particular point in the network is becoming congested, see if traffic can be redirected over other routes. You can change path costs to help redirect traffic.

It must be emphasized that when determining the throughput requirements of a network or line, peak loading situations must be considered as well as the average traffic volume. For example, if a network is designed for the average volume only, it may fail to handle peak loads well.

4.2 Guidelines for Configuring Ethernet LANs

If the end nodes on an Ethernet communicate only with each other, no routing node is required on the Ethernet. A routing node is needed where messages need to be routed off the Ethernet over other circuits such as DDCMP circuits.

If an Ethernet is operating with more than one area, and there are one or more level 1 routers on the Ethernet, then a level 2 routing node is required to transport messages between these areas.

Where possible, avoid dividing an Ethernet LAN into areas. The amount of routing overhead in Ethernet multiple area networks can become excessive. However, if the number of routers is excessive for a single area LAN, then it may be necessary and beneficial to divide the LAN into areas. See Sections 4.2.1 and 4.2.2.

When setting up any routing node on an Ethernet, you must follow the guidelines below to ensure proper node and routing operation. The guidelines deal with the number of routers and end nodes configured on an Ethernet.

1. The `MAXIMUM BROADCAST NONROUTERS` parameter defined on each Ethernet routing node should be greater than or equal to the total number of end nodes on all Ethernet circuits connected directly to the routing node within the same area (see Section 4.2.1).
2. The number of routers within the same area on an Ethernet should be kept to a minimum.
3. The `MAXIMUM BROADCAST ROUTERS` parameter defined on each Ethernet routing node should be greater than or equal to the total number of routing nodes that are on all Ethernet circuits connected directly to the routing node within the same area (see Section 4.2.2).
4. Any node to be used as a designated router should be reliable, having sufficient resources to handle the routing demands (see Section 4.2.3).

4.2.1 Number of End Nodes (Nonrouters) on an Ethernet LAN

The `MAXIMUM BROADCAST NONROUTERS` parameter defines the maximum number of end nodes for which the local routing node will store routing information. Define this parameter so that it is greater than or equal to the total number of end nodes on all Ethernet circuits connected to the local routing node. This total relates only to those end nodes within the same area as the routing node.

When two LANs are connected by a bridge, then the parameter's value should be greater than or equal to all the end nodes in the entire extended LAN. Again, if the network has multiple areas, the value should be geared to all end nodes in the same area.

NOTE

The maximum value of the MAXIMUM BROADCAST NONROUTERS parameter is 1022. If the total number of nonrouters in an extended LAN approaches or exceeds this maximum, then the extended LAN should be set up as a multiple area network. Then the limit will only apply on a per-area basis.

If the value assigned to the MAXIMUM BROADCAST NONROUTERS parameter is too small, the routing node's routing database will not be able to store information about all the end nodes on the Ethernet. Therefore, certain end nodes will not be reachable in one hop. (They may be reachable through other routers on the Ethernet which have a higher MAXIMUM BROADCAST NONROUTERS setting.)

Also, several conditions can occur that complicate troubleshooting and management. For example, if one node becomes unreachable, such as when the connecting line or system fails, this leaves an "opening" in the local routing node's routing database. Therefore, another node which had once been marked as unreachable because the parameter's limit had been exceeded will now become reachable.

As another example, if this parameter is set too low on one Ethernet router in the area, but is set large enough on another, then all routing traffic received by the former router and destined for nodes that exceed the parameter's value will be forwarded an extra hop through the second router.

Set the parameter to a value that is high enough to avoid the above mentioned problems. However, do not set the parameter any higher than is needed. (Some extra leeway can be given to accommodate future expansion of the network. This will save you from having to change the parameter every time an end node is added to the Ethernet.) If the value of the parameter is too high, this will introduce additional processing requirements on the local node for handling the routing overhead.

4.2.2 Number of Routers on an Ethernet LAN

Keep the number of routers within the same area on a single Ethernet to a minimum. The greater the number of routers, the greater will be the overhead for each router. The routers exchange routing information with each other, and so with a greater number of routers, there will be more of this routing traffic on the network. This traffic can become excessive, especially when the reachability status of nodes is fluctuating.

Although up to 32 routers are permitted in an area on an Ethernet, it is recommended that no more than 15 routers be used. Monitor routers to make sure they are not overloaded. (Use DECnet's Network Control Program to look at line counters. If there is an excess of "transit congestion losses" or buffer errors, the router may be overloaded.)

Configure more than one router to split a heavy routing load or to ensure that routing remains uninterrupted if one router goes down.

If your Ethernet requires numerous connections to nodes off the Ethernet or on other Ethernets, add one or more router servers or other dedicated systems that perform routing. Being dedicated solely to routing, these units can route without having to attend to other functions that normally occupy host-based routing nodes. This should increase the accessibility and reliability of your network.

In addition, the dedicated routers let you configure full-function host systems on the Ethernet as end nodes. Not having to perform routing, these end nodes can then dedicate themselves more fully to applications or other processing functions.

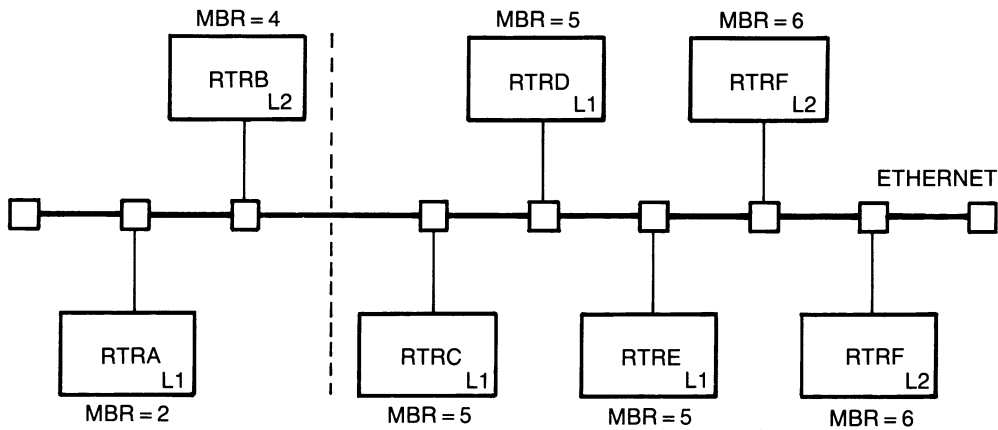
Define a routing node's `MAXIMUM BROADCAST ROUTERS` parameter so that it is greater than or equal to the total number of routers on the Ethernets to which the router is connected.

On a level 1 router in a multiple area network, the parameter must be greater than or equal to the total number of routers in the same area on all Ethernets to which the router is connected.

On level 2 routers, the parameter must be greater than or equal to the total of all level 1 and level 2 routers in the same area on the connected Ethernets, plus all the level 2 routers in all the other areas on the connected Ethernets.

To illustrate, assume the network shown in Figure 4-2 consists of two areas. Area 1 consists of one level 1 router and one level 2 router. Area 2 consists of three level 1 routers and 2 level 2 routers. The figure shows what value should be given to the `MAXIMUM BROADCAST ROUTERS` parameter (abbreviated `MBR` in the figure) on each of these routers.

For instance, on the level 1 router `RTRA` in area 1, the parameter is set to 2, since there are two routers in the same area on the Ethernet. However, on the level 2 router `RTRB`, the parameter is set to 4. This value incorporates the two routers in the same area plus the two level 2 routers in the other area.



LKG-0496

Figure 4-2: Defining the MAXIMUM BROADCAST ROUTER Parameter

In a multiple area network, the maximum number of routers that can be included in an area on an Ethernet is 32 minus the total number of level 2 routers in the other areas on the Ethernets connected to the router. (Again, the recommended maximum is 15 per area to prevent routing update traffic from congesting the Ethernet.)

In a single area network consisting of LANs connected by a bridge, the parameter's value should be greater than or equal to all the routing nodes in the entire extended LAN. If the network has multiple areas, the value should be geared to all level 1 routers in the same area plus all the level 2 routers in the entire network.

NOTE

If the total number of routers in an extended LAN approaches or exceeds the maximum of 32, then the extended LAN should be set up as a multiple area network. Then the maximum value will only apply to level 1 routers on a per-area basis.

If the value of the MAXIMUM BROADCAST ROUTERS parameter is too small, the local routing node may not be able to reach some other routers due to limits on the size of its routing database. The local routing node drops or excludes as many routers as necessary from its database to keep the total within the value of the MAXIMUM BROADCAST ROUTERS parameter. The routers that are excluded are those with the lowest routing priorities assigned to their Ethernet circuits (see Section 3.2.1.2). If two routers of equal priority exclude one another from their respective databases, this can lead to an unstable network.

Do not set the parameter any higher than needed. (Some extra leeway can be given to accommodate future expansion of the network. This will save you from having to change the parameter every time a router is added to the Ethernet.)

In some systems, if the parameter's value is too large, the memory overhead may also be large. Also, as said earlier, the more routers there are on the Ethernet, the greater amount of routing-control traffic will result.

On some routing systems there is a parameter that defines the maximum number of routers that the routing layer is to allow on a particular Ethernet circuit. This is the MAXIMUM ROUTERS parameter. (The MAXIMUM BROADCAST ROUTERS parameter is an executor node parameter, while the MAXIMUM ROUTERS parameter is a circuit parameter.) If the routing node has two Ethernet circuits, for example, you can use this parameter to specify the maximum number of routers for each circuit.

4.2.3 Potential Designated Routers

Any routing nodes that are likely to be chosen as designated routers should be reliable, and they should be able to handle all routing demands well. A node that is likely to be taken down often should not be allowed to be a designated router. To prevent it from becoming a designated router, the ROUTING PRIORITY parameter for its Ethernet circuit should be assigned a low value.

To ensure that the routing load is handled most efficiently, it is best to use dedicated routers as designated routers. Therefore, the ROUTING PRIORITY parameter should be set highest on any dedicated router(s) on the Ethernet.

4.3 Guidelines for Configuring Multiple Area Networks

Configuring a network that consists of multiple areas is more complex than configuring a network that, by default, consists of one area. The design of a multiple area network introduces a second, higher level of routing that links the areas. Designing a network for area routing involves awareness of certain network topological restrictions unique to area routing.

The area routing configuration guidelines presented below are based on these restrictions. The guidelines are intended to prevent problems such as loss of routing path, isolation of nodes, or incorrect routing of packets. These potential problems are discussed in Section 5.2.

When you configure a multiple area network, follow these guidelines:

- **Each node must belong in only one area.** This applies to all nodes in the network, Phase III nodes as well as Phase IV.
- **DECnet node numbers must be unique within an area.** However, they may be used again within another area. Thus, node identification within one area is independent of node identification in another area. If you combine two or more networks to form a multiple area network, you do not have to renumber all the nodes to maintain their unique identities. For example, node 211 in area 3 is identified as 3.211, while node 211 in area 16 is uniquely identified as 16.211.
- **Only a level 2 router can connect directly to a node in another area and thereby enable communication between the areas.** A level 1 router cannot have any circuits outside its own area.
- **Within a network, the level 2 routers must form a subnetwork; that is, they must be connected in such a way that they create a network of their own.** (See Figure 2-3.) There must be a level 2 routing path between any pair of level 2 routers across the network (without intervening level 1 routers, because level 1 routers cannot forward level 2 routing information).

When a level 2 router finds that it has no links to other level 2 routers, (either in its own area or in other areas), it automatically stops being a level 2 router. The router then functions as a level 1 router only. Conversely, when the level 2 router finds that links to other areas are active, it participates as a level 2 router.

- **Treat each area as though it were a separate network.** Each area must be physically intact and capable of running on its own. Within the area, there must be a level 2 path between any pair of level 2 routers.
- **Provide enough redundancy within each area and between areas to avoid having a single point of failure in the network.** For redundancy within an area, you could include more than one level 2 router and provide for alternate paths between nodes. This redundancy will prevent isolation of any one node or loss of the routing path within the area. For redundancy between areas, you could set up alternate paths between areas so that loss of a line does not disconnect any area from the rest of the network. Complete redundancy may not be feasible for small networks.
- The following restrictions apply to Phase III nodes:
 - Even though Phase III node addresses are not assigned area numbers, they are logically part of the area in which they reside.
 - All Phase III routing nodes are treated as level 1 routers. Therefore, do not link a Phase III routing node in one area with a node in another area. Such a connection could lead to area leakage, a problem described in Section 5.2.2.

- Never place a Phase III routing node in a path between two Phase IV nodes. A Phase III node cannot communicate with nodes in other areas or with nodes in the same area that have addresses greater than 255. To avoid placing a Phase III node on a path between two Phase IV nodes, place all Phase III nodes on the periphery of the network.
- **On each level 2 router, define the MAXIMUM AREA, AREA MAXIMUM HOPS, and AREA MAXIMUM COST parameters to suit your network configuration.** Follow the recommendations given in preceding sections of this chapter and those given in the DECnet system documentation for the local node.
- **Areas should reflect expected traffic flow.** That is, assign nodes to areas according to the frequency of their intercommunications or their likeness of logical functions. Although this guideline is not as strict as the previous ones, by following it you will greatly improve overall network performance. Routing within areas is less costly than routing between areas. Thus, areas should reflect not only physical proximity of nodes, but a combination of physical and logical proximity.

4.4 When and Where to Include Routers

The positioning of routers within a network is determined mainly by geography and the availability of communications lines. In general, you should use as few routers as possible. This will minimize the routing overhead (the overhead consisting of routing messages) and traffic on the network. However, provide for enough redundancy to maintain network reliability in the event of the failure of a router.

Include routers when network flexibility is required. End nodes are less flexible because they can have only one physical link.

Be aware that because routers are points of convergence for several paths, they are potential sources of congestion. Care should be taken not to overload a router with too many lines or too much network traffic. If routers are overloaded, the network's performance will degrade due to congestion packet losses. Each router's counters should be monitored to avoid this problem. Section 5.1.3 discusses how to tell whether a router is being overloaded and what you can do about it.

Ensure that the circuit bandwidths are adequate for the anticipated traffic. Remember that some routers can take advantage of paths of equal costs by using path splitting. However, make sure that all nodes in the network that will receive data transmitted over these paths support out-of-order packet caching.

4.4.1 Ethernet LANs

Routers can play an important role in Ethernet LANs, especially in connecting the Ethernet LAN to nodes or networks off the LAN. Where possible, full-function Ethernet host-based routing nodes should be converted to end nodes, leaving routing responsibilities to one or more server-based routers. This reduces system overhead and improves system performance on the converted full-function nodes.

You do not need routers on an Ethernet LAN unless connections to nodes off the LAN are necessary or unless the LAN is divided into areas (in which case one or more level 2 routers will be necessary in each area on the LAN). By including as few routing nodes as possible in an Ethernet LAN, you reduce the routing overhead in the network.

4.4.2 Multiple Area Networks

Each area of a multiple area network must have at least one level 2 router. All the level 2 routers must be able to reach each other without an intervening level 1 router.

When setting up multiple areas on an Ethernet, keep in mind that a bridge can be used to help isolate routing control traffic by restricting router multicast messages to the desired Ethernet. This implies that the areas on the extended LAN are separated physically by the bridge. If this is the case, then the bridge can be set up to filter the end node and level 1 Router Hello messages. This will limit the circulation of these messages to a single LAN (rather than having the messages propagate to the LAN on the other side of the bridge), increasing the available bandwidth on each LAN.

To take advantage of this feature, all the level 2 routers on the network must implement the latest DECnet routing software (version 2.1.0). Use the Remote Bridge Management Software (RBMS) to activate filtering of multicast messages. See the appropriate documentation for this product.

5

Routing-related Problems in Networks

This chapter discusses problems in networking, resulting from incorrect or inappropriate network configuration.

5.1 General Routing Problems


Several routing-related problems may occur in networks, each with several possible causes. The problems usually arise when inappropriate values are assigned to routing parameters. Trying to diagnose such problems can be very difficult, especially in large, complex networks. So it is very important to pay special attention to the guidelines and discussions given in the preceding chapters concerning the values of routing parameters.

5.1.1 Connectivity Problems

Most routing problems show up as connectivity problems: situations where a routing node is unable to reach or connect with one or more remote nodes. When such a problem occurs, verify that the routing node's routing parameters are defined appropriately for your network and that the remote node's DECnet software version is supported by the routing node:

Problem 1: Node Address Duplication

Connectivity problems may occur if two nodes have the same DECnet node address. In Figure 5-1, nodes X and W have the same address. When node X tries to communicate with node Y, node Y responds by communicating over the shortest path to what it thinks is node X, but what it thinks is node X is really node W. Therefore, node X never receives messages back from node Y and assumes node Y is not reachable.



LKG-0498

Figure 5-1: Nodes with Equal Addresses

Remedy: Make sure all nodes of the network have unique DECnet node addresses. Guidelines for defining node addresses are given in Section 4.1.1.

Problem 2: MAXIMUM ADDRESS Parameter Too Small

If a routing node cannot reach a node, the problem may be that the MAXIMUM ADDRESS parameter on the routing node is less than the DECnet node number of the unreachable node. The value of the MAXIMUM ADDRESS parameter should be greater than or equal to the largest node number in the network. In multiple area networks, it should be greater than or equal to the largest node number in the local area.

Similarly, the problem may be that Phase III and Phase IV nodes coexisting in the same network are not configured properly. If the local routing node is a Phase III node, then it will not be able to reach any Phase IV nodes with addresses exceeding 255. If the local routing node is a Phase IV node, then perhaps there is a Phase III routing node on the path to the destination. The Phase III routing node cannot route to any nodes whose addresses exceed 255.

If event logging is enabled, the routing node logs a "Node out of range packet loss" event message when discarding a packet destined for a node whose address

exceeds the MAXIMUM ADDRESS parameter. The “Partial routing update loss” message is logged when the router receives a routing message containing node addresses that exceed the MAXIMUM ADDRESS parameter.

If the adjacent node is connected to the routing node by a nonbroadcast circuit and the adjacent node’s address exceeds the routing node’s MAXIMUM ADDRESS parameter, the circuit may be taken out of service. If event logging is enabled, the routing node logs a “Circuit down” event message.

Remedy: If the value assigned to this parameter is too small, increase it to the proper size. If the number of nodes in your network exceeds the maximum value that can be given to this parameter (255 in Phase III networks, 1023 in Phase IV networks), then convert your network into a multiple area network (see Section A.1).

If Phase III and Phase IV nodes coexist in the same area, check that all Phase IV nodes with which Phase III nodes communicate have DECnet node numbers less than 255. Also, no Phase III nodes should be in a path between two Phase IV nodes.

Guidelines for defining the MAXIMUM ADDRESS parameter are given in Section 4.1.2.

Problem 3: MAXIMUM BROADCAST NONROUTERS Parameter Too Small

In Ethernet networks, a routing node may not be able to reach an end node on the same Ethernet LAN because the value specified for the MAXIMUM BROADCAST NONROUTERS parameter is not large enough to include all the end nodes on the Ethernet.

If event logging is enabled, the routing node logs an “Adjacency rejected” event message.

Remedy: If the value of this parameter is too small, increase it. If your network is a single area network and the number of end nodes in your network exceeds the maximum value that can be assigned to this parameter (1022), consider converting your network into a multiple area network. In multiple area networks, this parameter applies to the number of nodes in the local area only.

Guidelines for defining the MAXIMUM BROADCAST NONROUTERS parameter are given in Section 4.2.1.

Problem 4: MAXIMUM BROADCAST NONROUTERS Parameter Too Large

Another reason why a routing node cannot reach certain end nodes may be that the value of the MAXIMUM BROADCAST NONROUTERS parameter exceeds the MAXIMUM ADDRESS parameter. (The router cannot address nodes with DECnet node numbers exceeding the MAXIMUM ADDRESS parameter.)

Remedy: Either decrease the value of the MAXIMUM BROADCAST NONROUTERS parameter so that it does not exceed the MAXIMUM ADDRESS parameter, or increase the MAXIMUM ADDRESS parameter so that it is greater than or equal to the MAXIMUM BROADCAST NONROUTERS parameter.

Problem 5: MAXIMUM BROADCAST ROUTERS Parameter Too Small

In Ethernet networks, a routing node may not be able to reach another routing node on the same Ethernet LAN because the value of the MAXIMUM BROADCAST ROUTERS parameter is too small to accommodate all the routers on the Ethernet.

When this parameter is too small, the local routing node excludes as many routers as necessary from its routing database to keep the total within the maximum value of the MAXIMUM BROADCAST ROUTERS parameter. The routers that are excluded are those with the lowest routing priorities assigned to their Ethernet circuits.

If event logging is enabled, the routing node logs an “Adjacency rejected” event message.

Remedy: If the value of the parameter is too small, increase it to the proper size. If your network is a single area network and the number of routers in it exceeds the maximum value that can be assigned to this parameter (32), then consider converting your network into a multiple area network. Some of the routers will have to be assigned to remote areas as level 1 routers.

Guidelines for defining the MAXIMUM BROADCAST ROUTERS parameter are given in Section 4.2.2.

Problem 6: MAXIMUM PATH COST or MAXIMUM HOPS Parameter Too Small

A routing node may not be able to reach a node because the path with the lowest cost to that node exceeds the value of the MAXIMUM PATH COST parameter, or the path with the shortest distance to that node exceeds the value of the MAXIMUM HOPS parameter.

If event logging is enabled, the routing node logs a “Node unreachable packet loss” event message.

Remedy: Make sure the values of these parameters are set large enough so that all nodes with which the routing node is to communicate will be reachable.

Guidelines for defining these parameters are given in Section 4.1.5.1.

Problem 7: Unsupported Software on Remote Node

The routing node may be trying to communicate with a node that is running unsupported DECnet software.

If event logging is enabled, and the unreachable node is adjacent to the routing node, the routing node logs a “Circuit down” message along with the reason for the event: “Version skew.”

Remedy: Check that the DECnet software on the remote node is supported by the local routing node. For example, Phase IV nodes only support connections to Phase IV or Phase III nodes.

5.1.2 Data Link Problems

Communications problems may arise because of mismatches on the data-link level. For example, there may be a mismatch between line speeds. Either (1) the line speed specified in a node's configuration database does not match the actual speed of the line device, or (2) the line speeds for a line connecting two adjacent nodes, as specified in the configuration databases of the two nodes, do not match.

If the routing parameters and data link characteristics have been verified, examine the event messages generated by the router. Also, examine the counters on the circuit between the router and the node which cannot communicate with the router.

To follow are two other problems that commonly arise because of mismatches on the data-link level:

Problem 1: Level 1 Router Connected to Another Area

If the local router is a level 1 router and it is connected directly to a node in another area, then the local router will not be able to establish a link with the connected node.

Remedy: Level 1 routers only communicate with nodes within their own area. Change the address of the attached node so that it is in the same area as the router, or convert the router to a level 2 router, if appropriate.

Problem 2: Verification Reject

If the local router and/or the attached node require verification and their receive and transmit passwords do not match, then the nodes will not communicate. If event logging is enabled, the routing node logs a "Verification reject" event message.

Remedy: To see if this is the problem, disable verification on both nodes and retry the initialization between them.

If the nodes do initialize successfully, then the passwords do not match. Correct the passwords and enable verification again. Remember that the transmit password on one node must match the receive password on the other.

If after examining the possible problems and causes suggested in this and the preceding section, you still do not find the cause, perform data link tests (circuit loop-back tests) to verify connectivity. Also, execute the NCP SHOW NODE command at the local router, specifying the node to which the connection cannot be made. If the SHOW NODE command does not work, the problem lies in the router.

5.1.3 Performance or Lost Packet Problems

If a routing node's performance is degrading or if counters reveal that packets are being discarded, the problem may be that traffic demands are too much for the router. As discussed below, another router may be needed to off-load some of the traffic from the first router.

However, the problem may be because of the way buffers are allocated, causing the router to perform below its potential. The buffer sizes throughout the network may be inconsistent, or the buffer size on the routing node may be too small.

Consider the following possibilities:

Problem 1: Inconsistent Buffer Sizes Throughout the Network

Communications problems can occur if the size of network buffers is not consistent throughout the network.

If the local router's network buffer size is smaller than that of other nodes, then the router may be receiving packets that are bigger than its buffer size. The router discards those packets.

If a router is on a path between two nodes whose buffer sizes exceed that of the router, packets exchanged between the two nodes will be dropped by the router.

If the buffer size of the next node on the path to the destination is smaller than that of the routing node, or smaller than the packet being forwarded by the routing node, the routing node will discard the packet.

If event logging is enabled on the router, it logs an "Oversized packet loss" event message when packets are discarded.

A "Circuit down" message with the reason "Adjacent node block size too small" indicates that the buffer size on the node adjacent to the router is less than the minimum acceptable buffer size. The minimum is 236 bytes. However, for networks of greater than 14 nodes, the minimum acceptable buffer size increases with the size of the network. Use the following formula to compute the minimum:

Minimum Buffer Size = 2 x (MAXIMUM ADDRESS) 9

Remedy: Check that buffer sizes are consistent throughout the network. Guidelines for defining buffer sizes are given in Section 4.1.3.

Problem 2: Insufficient Buffers Allocated

If routing systems that dynamically allocate buffers do not allocate enough buffers to handle communications demands, communications problems will occur.

Remedy: Check that enough buffers are being allocated on the local routing node.

On systems where buffer allocation is fixed, if the buffer size is excessively large, fewer buffers can be allocated (because of memory limitations on the node). With fewer buffers, there may not be enough buffers to handle all the incoming data.

To diagnose this problem, execute the NCP SHOW CIRCUIT *circuit-id* COUNTERS command on the routing node. On Ethernet circuits, look at the "User buffers unavailable" counter. On point-to-point circuits, look at the "Remote buffer errors" and "Local buffer errors" counters.

If any of these have high counts, then the number of buffers in the system may be insufficient. Make sure the buffer allocation is appropriate for your network application (see Section 4.1.3).

On either Ethernet or point-to-point circuits, look at the "Transit congestion losses" counter. If it is excessive, the problem may be that the number of buffers (packets) queued to a circuit exceeds its limit. Routing packets will be discarded because of this. This limit is specified by the ROUTING QUEUE THRESHOLD parameter. On some systems, this parameter is determined automatically. On others, the network or system manager can adjust its value to tune the operation of the router. The value should be based on the number of buffers allocated on the router.

The congestion losses may be the result of an insufficient number of buffers to handle the load (perhaps because the buffer size is too large, as discussed above). It may also be that a high-speed line brings too much traffic to the router, as discussed below.

Problem 3: End-to-end Communications or System Performance Degradation

If the performance for end-to-end communications is degrading, or if the router is a full-function routing node and system performance is degrading, it may be because of insufficient buffer allocation or excessive routing traffic.

On many systems, routing has higher priority than end-to-end communications or system processes. So when the load on the router becomes excessive, certain end-to-end communications or system processes will be deferred in favor of routing.

Remedy: For some critical full-function routing nodes in large networks, it may be necessary to guarantee that user processes running on the node never interfere with the memory requirements of the network software. There should always be enough memory to support buffers for network processes such as routing; otherwise, system performance will be very slow. System processes will continually be put on hold while they wait for available buffer space.

If the lack of memory makes the network software on the node unable to allocate a buffer fast enough to receive data from a communications line, the line may be considered unusable by another node in the network. When this happens, the network attempts to adaptively reconfigure its routing databases, resulting in network traffic consisting of configuration update messages.

If the node with memory problems should be close to failing, without failing completely, it may alternate between working and not working, causing the network to repeatedly reconfigure itself. Ultimately, these reconfigurations will degrade the performance of the entire network.

For more information on tradeoffs when using full-function nodes as routers, refer to the appropriate system DECnet documentation for the node.

Problem 4: Routing Performance Degradation

Under certain circumstances, a router's performance may degrade. The performance problem may appear on one or all of the connected circuits.

Remedy: When a router's performance appears poor, check if it is just one or two lines that are overloading the router or if it is the aggregate demand of all the router's lines. Line congestion can be diagnosed by looking at the "Transit congestion loss" counter for each circuit, displayed by executing the NCP SHOW CIRCUIT command.

Compare this counter with the "Transit packets sent" counter for the circuit. If too many packets are queued for transmission over a circuit, congestion will occur. You can tell the degree of congestion by comparing these counters. If the number of congestion losses is large relative to the number of packets sent, serious congestion has occurred over that circuit. This is possibly due to an overloaded routing node or an overloaded circuit.

An overloaded circuit can occur, for example, when two LANs are connected by a single synchronous or asynchronous line. In this configuration, if the communications between the two LANs delivers more packets than can physically be sent on the serial line, then transit congestion will occur. One way to tell if this is occurring is to examine the circuit counters for the line in question. Follow these steps:

1. Zero the counters and let them accumulate counts for a while.
2. Record the counts shown by the "Bytes sent" and "Seconds since last zeroed" counters.
3. Divide the number of bytes sent by the number of seconds since the counters were zeroed.
4. Multiply this result by 8 for synchronous lines, or by 10 for asynchronous lines.
5. Compare the result with the speed of the line.

If the results indicate an appreciable percentage of the line is being used (that is, over 70 to 80%), then the transit loss is probably due to an overloaded line. (Also note that since half-duplex lines are shared by both ends of the communications line, the percent-utilization figures should be computed for both ends of the line and their sum should be compared to the line speed.)

If the problem is an overloaded line, then either a higher speed line or multiple path-splitting lines (if possible) should be used.

If only one or two lines are affected, meaning that the other lines are handling their loads sufficiently, then you can remove the congestion by adding another line. If you can set up two paths of equal costs and implement path splitting, this will lessen the traffic on any one line.

If all lines seem to be affected, then the routing traffic exceeds the capacity of the router. You can tell whether the router is overloaded by how slowly the router responds to network management commands executed on the router. On many routers, you can use a monitor utility to see how busy the CPU is.

If the aggregate throughput demand on the routers' circuits exceeds the router's capacity, performance will degrade. If this is the case, you may have several options, such as:

- Upgrade the router to increase its capacity.
- Add another router to share routing demands.
- Redirect some of the routing load to other routers and paths (for example, by increasing circuit costs so that other paths are used).

If your node is an Ethernet router and it seems to be overloaded by network traffic, it may be because it is the designated router. The designated router takes on an extra load because it must handle routing for end nodes wishing to communicate with nodes off the LAN (and for first communication attempts with other end nodes on the LAN). The designated router is the router with the highest routing priority assigned to its circuit. Make sure any potential designated router has sufficient capacity to handle the routing demands.

To determine how well a router is performing, compare the amount of traffic it is routing to the router's rated throughput. The router's **rated throughput** is usually stated in the Software Product Description.

To calculate the amount of traffic, do the following:

1. Zero all circuit counters by using the NCP ZERO KNOWN CIRCUIT COUNTERS command.
2. After a while, execute the NCP SHOW ACTIVE CIRCUIT COUNTERS command and check the "Transit packets received" counter for each circuit.
3. Take the total of all transit packets received on all circuits and divide by the number of seconds lapsed since the counters were zeroed.

This will give you the number of routing packets processed by the router per second (its **real throughput**).

Set up event logging and counter timers to force logging of the counters at regular intervals. You should check the event-logging messages and counters closely at all times to avoid potential overload problems. By checking counters at timed intervals, you can then calculate the average packets/per second count for each circuit.

(Remember to take several long samples, as traffic may be intermittent: if your timed interval happens to be during a lull in traffic, your calculations will not reflect the true extent of the traffic load.)

Problem 5: Path-Splitting Problems

Where path splitting has been set up, performance problems can occur if the paths used for splitting the traffic are dissimilar. They should not have lines of widely diverging speeds, nor should one path have many more hops than another; otherwise, this situation could force an excessive amount of out-of-order packet caching on destination nodes. This may overload these destination nodes and cause performance problems.

Remedy: Where path splitting is implemented, check that the paths used for this function are similar.

Problem 6: Excessive Network Delay Due to Path Length

If the **network delay** is excessive, the reason may be that the paths being used are too long. If paths have numerous hops, the network delay becomes large because of the cumulative delay incurred for routing at each of the intervening nodes on the path.

Remedy: Try to keep path lengths to a minimum where possible. Section 4.1.5 discusses when to configure long or short paths.

5.2 Problems in Multiple Area Networks

This section discusses some typical problems specific to multiple area networks and offers suggestions to help you resolve those problems.

Problem 1: Problems Reaching Other Areas

If a routing node cannot reach another area, it may be because the values specified for the **MAXIMUM AREA**, **AREA MAXIMUM COST**, or **AREA MAXIMUM HOPS** parameters are too low. For example, if the value of the **MAXIMUM AREA** parameter is 10 and there are 12 areas, then two areas will be unreachable.

Remedy: The remedy for this problem is similar to those discussed for the **MAXIMUM ADDRESS**, **MAXIMUM PATH COST**, and **MAXIMUM HOPS** parameters discussed in Section 5.1.1.

Problem 2: Area Leakage

Communications can occur inadvertently across area boundaries. This can happen when a Phase III node is connected to a node outside its own area. For example, node **DAMION** in Figure 5-2 is a Phase III node with a link between areas 3 and 4. Because of the way a Phase III link handles an incoming node address (it drops the area number), **DAMION** will build a routing database containing nodes **ABLE**, **CANE**, and **EAGLE**, including their node addresses but not their area numbers.

DAMION will send routing updates about ABLE and CANE across the area boundary to EAGLE. Because the node addresses of ABLE and CANE sent to EAGLE will not include the area number, EAGLE will think that ABLE and CANE are in its own area. Likewise, ABLE and CANE will think EAGLE is in their own area. In this way, nodes ABLE and CANE will be able to communicate with EAGLE, and vice versa. This kind of configuration actually works in some cases. However, when it does not work, it can be very difficult to diagnose and solve the problem. Therefore, do not configure your network in this way.

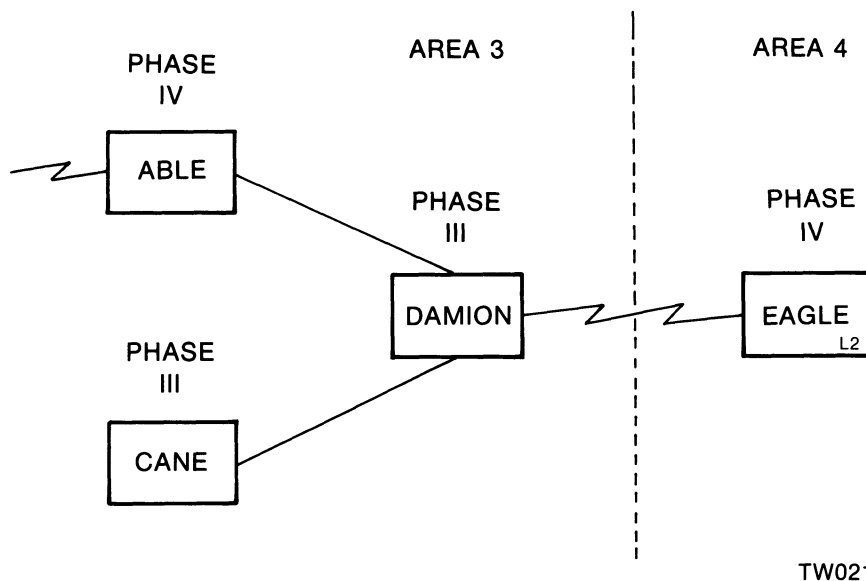


Figure 5-2: Area Leakage

Because of the possibilities of area leakage, do not use areas to enforce protection. Although at first glance the concept of areas seems to offer penetration protection across areas, the DNA architecture does not specify protection guarantees for areas. Therefore, do not assume such guarantees when you set up your multiple area network.

Remedy: If a Phase III node must communicate with a node in another area, either upgrade the Phase III node to Phase IV, or configure the local area so that the Phase III node can communicate through intervening Phase IV nodes (including at least one level 2 router) that can reach the remote area.

DECnet Phase IV provides a way for you to prevent the interarea leakage problem described above. Before a link can be established between a Phase III and Phase IV node, the Phase IV node requires a node level password (transmit password) from

the Phase III node. If the Phase IV node does not receive the required password, the link will not be established (the circuit will not come up). The event-logging facility records this as an error message, indicating that a password is required or is mismatched. To properly use this feature to help configure Phase III routing nodes, assign a separate password (used on Phase III nodes) for each area.

This helps you locate interarea links between Phase III and Phase IV nodes. You can learn which Phase III nodes have not been initialized. You can use this information to prevent Phase III nodes from linking to nodes outside their own areas, or to identify which Phase III nodes need to have the transmit password set. Note that this technique will not locate Phase III nodes that are improperly linked to nodes in other areas if the Phase III nodes were configured using node level passwords prior to area conversion.

When you are sure your network is configured properly, define transmit passwords for all Phase III nodes and receive passwords for all Phase IV nodes. The receive password defined on a Phase IV node should match the transmit password of any Phase III node connected to it.

Problem 3: Single Points of Failure

If only one or two level 2 routers are configured in an area, then if those routers should become unavailable for some reason, that area will become isolated from the rest of the network.

Remedy: You should configure several level 2 routers in the same area to provide redundancy. In this way, when one of the routers is unavailable, routing to and from other areas can continue through the other level 2 router(s). Provide for alternate paths between nodes in the same area so that the loss of one path does not prevent the flow of level 2 routing traffic through the area.

Problem 4: Potential Problems with Area Routing on an Ethernet

You can set up multiple areas on an Ethernet; all the nodes on a single Ethernet do not have to be in the same area. However, try avoiding such configurations where possible; otherwise, the message overhead for routing nodes (especially level 2 routers) on the Ethernet increases.

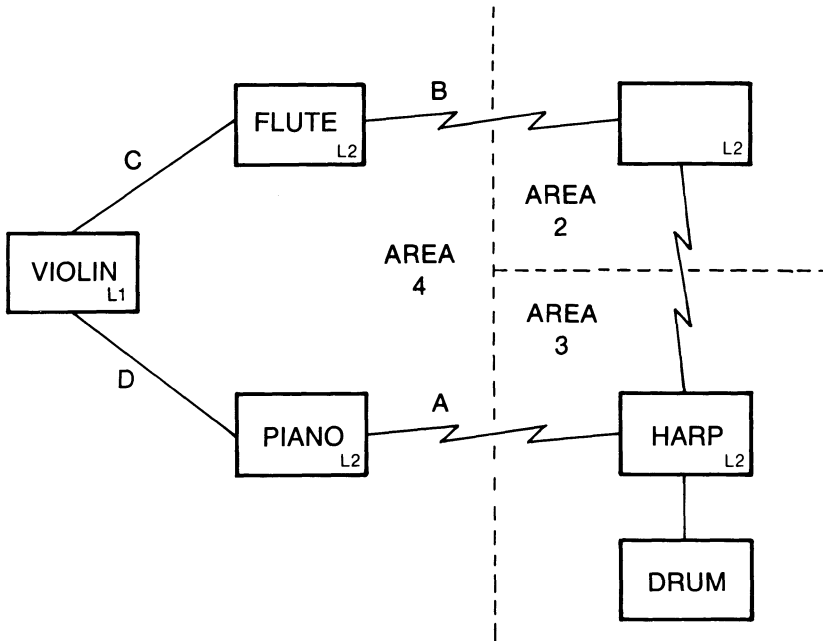
When two end nodes or level 1 routing nodes on an Ethernet are configured in different areas, the nodes do not communicate directly with each other. Each node first communicates with a level 2 router in its own area, which sends the message to a level 2 router in the other area. That level 2 router then transmits the message to the destination node.

Remedy: Avoid dividing an Ethernet LAN into areas, except when the number of routing or end nodes exceeds their respective allowable limits. If you do implement area routing in an Ethernet LAN, make sure all level 2 routers have sufficient processing power and memory to handle routing traffic.

Problem 5: Improper Configuration of Routing Nodes Within an Area

If a multiple area network is configured improperly, network traffic may be routed incorrectly or lost. Figure 5-3 shows an improperly configured area in which a node can be isolated within an area. This problem is called **area partitioning**, where an area is vulnerable to partitioning as a result of failure of a single line.

Assuming each link cost is 1, then if link C or D is down, the following can happen: Assume we want data sent from node DRUM in area 3 to node FLUTE in area 4. The level 2 router in area 3, HARP, will determine the least cost path to a level 2 router in area 4, which is the path from DRUM to HARP to PIANO. But if link C or D is down, then FLUTE will be isolated. On an initial attempt to connect to FLUTE, DRUM will receive a "Node unreachable" message. If link C or D breaks after DRUM and FLUTE have established a link, a timeout will occur.



LKG-0566

Figure 5-3: Improper Router Configuration Within an Area

Another problem can occur with the configuration shown in Figure 5-3. If link D is down and node FLUTE wishes to create a link to node HARP, node FLUTE will choose a path through area 2 and then to HARP. However, when HARP tries to send a return message, the same problem will occur as described above.

Remedy: To prevent these problems when configuring your area network, remember to treat each area as a separate network. Where possible, each node should have alternate paths to other nodes, and all level 2 routers in an area should be connected by level 2 paths. In Figure 5-3, the network in area 4 is poorly configured because only one path can be used by each node to reach another node in its area. Also, a level 1 router is on the path between two level 2 routers (node VIO-LIN is between node FLUTE and PIANO). One way to fix the above problems is to install a link between FLUTE and PIANO.

Problem 6: Phase III and Phase IV Coexistence

Though DECnet Phase IV networks do support coexistence of Phase IV and Phase III nodes, complications can arise if the nodes are not configured properly. Problem 2 has shown what happens when you set up a Phase III node with a link in to another area (Figure 5-2). The following describes the problems that occur when you place a Phase III routing node on a path between two Phase IV nodes.


To understand the cause of these problems, you should understand how a Phase III node or link behaves within a network that supports area routing. Phase III nodes do not know anything about areas. Therefore, when data to or from another area is routed through a Phase III node, the consequences can be unpredictable.

Whenever a Phase III node links with a Phase IV node in an area-based network, the circuit is initialized with the Phase III protocol and all references to the area number of nodes are dropped. This affects the network in two ways:

1. If the traffic is going from the Phase IV node with area routing to the Phase III node, the area designation is dropped. For example, if the address for node 5.167 is passed to a Phase III node, the address will become node 167 and the Phase III node will pass it on that way.
2. If the traffic is going from a Phase III node to a Phase IV node, any intervening routing nodes add their own area number to the address of the destination node stored in the routing packet. For example, if a Phase III routing node in area 5 forwards a packet destined for a node with the address 8.167, the Phase III node drops the area designation and passes the packet to the nearest level 2 router in area 5. The destination address on the packet is now 167 instead of 8.167. The level 2 router will then add its own area number to the destination address, making it 5.167. Thus, the packet will never reach its intended destination. Of course, if the destination node were in the same area as the level 2 router, the traffic would reach its proper destination.

Referring to Figure 5-4, the following shows what may happen when a Phase III node is placed between two Phase IV nodes: If node SALES (node address 4.45) wants to communicate with node ACCNT (node address 4.88), data first goes to the Phase III node CHECK, which has a node address of 22 (no area number for Phase III nodes). Node CHECK drops the area number from the destination address of the data; that is, ACCNT's node address becomes 88. Then CHECK routes the data to ACCNT. ACCNT adds its own area number to the node address. The destination address becomes 4.88. Because the entire path is within the same area, no problem results.

However, if node WAREH (node address 2.72) wants to send data to node SALES (4.45), the following happens: Because these two nodes are in different areas, WAREH sends the data to the nearest level 2 router in the destination area, which is node ACCNT. In turn, ACCNT routes the data to the Phase III node, CHECK, where all references to area addresses are dropped. CHECK sends the data to SALES with no problem. But the problem arises when SALES tries sending a message back to WAREH: to send the message to WAREH, SALES must locate the nearest level 2 router. SALES depends on node CHECK for information about level 2 routers. But because CHECK is a PHASE III node, no such information is available. Therefore, SALES will not forward the message to CHECK. Eventually, node WAREH will time out while waiting for a response from SALES.



TW023

Figure 5-4: Improper Configuration of Phase III and Phase IV Nodes

Remedy: To avoid this problem, switch the locations of nodes SALES and CHECK so that CHECK is not between the two Phase IV nodes.

As mentioned previously, try to upgrade all Phase III nodes to Phase IV where possible. If you must have Phase III nodes, be careful how you configure them.



A

Configuring a Multiple Area Network

This appendix discusses the steps to follow when configuring a multiple area network. More guidelines are discussed in Section 4.3. Section A.1 discusses the steps for converting a conventional single area network into a multiple area network. Section A.2 discusses the steps for putting together a multiple area network from start to finish.

A.1 Converting a Standard Network into a Multiple Area Network

Converting an existing single area network to a multiple area network requires careful planning. Because network addresses of existing nodes will change, some nodes may be unreachable while the conversion is underway. Follow these steps to keep this disruption to a minimum:

1. Decide on the new network topology, using all the guidelines in Section 4.3. Decide which nodes should be level 2 routers, level 1 routers, and end nodes.
2. If the new design entails repositioning some nodes, make the required changes before you begin converting node addresses. For example, the redesign may require reconnecting a Phase III node so that it is not in a path between two Phase IV nodes.
3. Create new node databases reflecting the new layout.

The major changes to make to the permanent database of each node are:

- Establish area support on each node, starting with the level 2 routers, support by redefining each node's address to correctly reflect its area.

- Define the level 2 routers. Most DECnet Phase IV routing nodes can be upgraded to a level 2 router. Level 1 and level 2 routers are distinguished by the TYPE parameter in the node's executor database. The node type of level 1 routers is ROUTING IV. The node type of level 2 routers is AREA. Refer to the appropriate documentation to find out how to modify these parameters.
 - On each node, redefine the remote node names, assigning the correct area numbers for the network.
 - Define transmit passwords for Phase III nodes connected to Phase IV nodes. The transmit password on the Phase III node must match the receive password for the adjacent Phase IV node.
4. Shut down the network and bring it up again with the new database.

This is the only part of the conversion process that benefits from real-time cooperation among the nodes in your network. If you can take down an entire network or area to upgrade each node, do so. If you cannot shut down all nodes simultaneously, such as in an environment that cannot tolerate massive disruption of the network, some nodes will not be able to interconnect until all nodes have been brought up with the new database.

Avoid using area number 1 for any area in the revised network, since node address duplications might occur during the transition period.

At the same time, each node in the network should shut down DECnet. It is not necessary to shut down the operating system.

After all nodes have shut down DECnet (or after an agreed-upon interval during which all nodes should have shut down DECnet), restart DECnet on all the nodes.

Note that if your node is on an Ethernet to which applications other than DECnet are connected, such as the Local Area Transport (LAT) protocol used for communication with terminal servers, these applications should also be shut down along with DECnet, and then restarted after DECnet is restarted. This step is necessary because DECnet will be changing the Ethernet physical address of your node to reflect the new executor node address.

5. Try communicating with other nodes in the network.

First try to communicate with nodes in the same area. If no problems result, try to communicate with nodes in other areas. If you are unable to communicate with another node, use NCP SHOW commands to try to diagnose the problem. For example, issue the following commands:

```
NCP SHOW NODE node-id STATUS
```

Assuming the node identified by *node-id* is the node with which you were unable to communicate, this command will show you the next node in the path to that node and the name of the circuit used.

NCP SHOW CIRCUIT *Ether-circuit-id*

This shows all adjacent nodes on the identified Ethernet circuit. Make sure all the nodes you expect to see are there (and in the correct area).

NCP SHOW CIRCUIT *Ether-circuit-id* CHARACTERISTICS

This shows the designated router on the identified circuit.

Try to determine the path to the node with which you were unable to communicate. Use the NCP SHOW NODE command described above to find out the next node along the path. Once you know that node, issue the command remotely at that node to find the next node along the path. Continue in this way until you come to the destination node. By "walking" down the path to the destination, you may be able to find the point of failure.

For example, assume the destination node with which you were unable to communicate is node BACH. First, issue the following NCP command at your local node:

```
NCP SHOW NODE BACH STATUS
```

Assume the SHOW NODE display reveals that the next node on the path to node BACH is node MAHLER; then issue the following NCP command from your local node:

```
TELL MAHLER SHOW NODE BACH STATUS
```

Continue in this way until the SHOW NODE command displays node BACH as the next node on the path.

NOTE

If a node along the path does not have node name BACH defined in its node database, the NCP TELL command will fail. You may have to use the DECnet node address instead of the DECnet node name.

Depending on the size of the network and the care with which the conversion is done, there may be a period requiring that you debug the network to ensure that the desired level of connectivity has been achieved. You can simplify debugging if you can run each area separately for a while before connecting them together. You can do this by turning off the circuits between level 2 routers in different areas.

Once you are confident that an area is operating to your satisfaction, you can turn on the circuits joining the area to its neighboring areas. Of course, while the interarea circuits are off, nodes will not be able to access other areas. View this as a tradeoff to reduce the number of variables during conversion.

A.2 Designing Multiple Area Networks

Use the following guidelines when designing multiple area networks.

1. Decide what areas you need. Base your decision on such factors as the number of nodes, geographical location, expected traffic patterns, and availability of lines.
2. Design the level 2 routers, area by area, into a level 2 subnetwork. Then, in each area, add the level 1 routers. Finally, add the end nodes required to complete each area. Do not connect the areas together until you have configured all areas in the network, as explained previously, and they are running and stable.
3. After configuring all level 2 routers as a “backbone” network, tie all of the areas to it. Remember that the level 2 routers must form a subnetwork so that each level 2 router directly connects to another level 2 router. Level 1 routers or Phase III nodes cannot form a part of the path between two level 2 routers. To ensure reliable area routing, provide more than one path between level 2 routers within the same area.
4. Make sure each node belongs to one area only.
5. Only level 2 routers can be one hop away from a node in another area. That is, only level 2 routers can have a link connecting any two areas.
6. Provide enough redundancy in each area to handle failures of links and nodes. As mentioned earlier, it is advisable to set up several level 2 routers in each area to prevent disruption of area routing when a single level 2 router goes down.
7. Treat all Phase III routing nodes as level 1 routing nodes. All Phase III nodes (both routing and nonrouting) can communicate only with nodes within the same area. Place all Phase III nodes on the periphery in each area so that no Phase III node has a direct connection to a node in another area. Do not place a Phase III node on a path between two Phase IV nodes.

Glossary

Adaptive routing

A form of routing in which messages are forwarded through the network along the most cost-effective path currently available, and are automatically rerouted if a circuit becomes disabled. In DECnet, a routing node's path choice depends on the current information stored in its routing database, which is updated regularly and instantaneously as conditions change, to reflect the latest condition of the network. See also **Routing**.

Adjacency

The node and circuit (in combination) connected to the local node.

Adjacent node

A node next to the local node, which is attached directly to the local node by a single physical line.

Aged packet

A packet that has exceeded the maximum number of visits allowed packets in the network. See **Maximum visits**.

Aggregate throughput

See **Throughput, network**.

Area

A group of nodes within a network.

Area router

A level 2 router in a multiple area network that performs routing between areas.

Area routing

The forwarding of packets from one area within a network to another area using level 2 routers. Nodes in a network are grouped into areas for routing purposes. Routing in a multiple area network is hierarchical, with one level of routing within an area (level 1 routing) and a second, higher level of routing between areas (level 2 routing).

Asynchronous transmission

A mode of transmission in which the time intervals between character transmissions differ. Each character is surrounded by a start bit and stop bit(s) to allow the receiving device to recognize the beginning and end of each character. (Also called "start-stop transmission.") Asynchronous transmission most commonly occurs over terminal lines and lines connecting other inexpensive devices (such as personal computers). Contrast with **Synchronous transmission**.

Availability

The proportion of time a specific piece of equipment, system, or network is usable, compared to the total time it is expected to be.

Bandwidth

The bandwidth of a transmission medium is the range between the highest and lowest frequencies at which signals can be passed through it without the signal being distorted beyond recognition by the receiver. The bandwidth reflects the medium's information-carrying capability. Usually expressed in terms of its signaling rate in Hertz or bits per second.

Block

Any contiguous unit of user information that is grouped together for transmission, such as the user data within a packet, excluding the protocol overhead.

Bottleneck

A point in the network where traffic is delayed or blocked. Usually a saturated device becomes a bottleneck. Bottlenecks are the limiting factors in network performance. See also **Congestion**.

Bps

Bits per second.

Bridge

A device that expands the extent of a LAN by connecting it to another LAN or physical link. For example, Digital Equipment Corporation's LAN Bridge 100 manages inter-LAN traffic and selectively forwards packets to keep local traffic local. Only data destined for different LANs passes through the bridge and continues on to the appropriate remote destination. Thus, the bridge improves performance by reducing traffic between LANs.

Broadcast

In networks, the capability of sending a single message simultaneously to many nodes, just as television programs are sent to multiple television sets.

Broadcast addressing

This is a type of multicast addressing in which all nodes receive a message simultaneously.

Broadcast circuit

A circuit connecting multiple nodes and capable of transmitting messages to multiple receivers.

Buffer

A device or an area of memory used for temporary storage to compensate for a difference in rate of data flow or in time of occurrence of events, when transmitting data from one device to another. Buffers are used on routing nodes to temporarily store data that is to be forwarded from one node to another.

Buffering level

In network communications, this refers to the number of buffers provided at a time by the network software to handle data. Single-buffering tends to be less efficient than multibuffering but uses less memory on the local system. Multibuffering provides better performance if sufficient memory is available. With multibuffering, a network can send or process several buffers of data in quick succession.

Bus

1. A LAN topology where all nodes connect to a single transmission medium so that all nodes are equal and all nodes hear all transmissions on the medium. Bus topologies are reliable because failure of any node does not affect the ability of any other nodes to transmit and receive.
2. Digital-specific: A flat, flexible cable consisting of many transmission lines or wires and used to interconnect computer system components to provide communication paths for addresses, data, and control information.

Cache

1. The process of storing blocks in memory for future use; used to minimize physical transfer of data between mass storage devices and memory.
2. A very fast memory used in combination with slower, large capacity memories.

Carrier sense

A signal provided by the Physical layer to indicate that one or more stations (nodes) are currently transmitting on the Ethernet channel.

Carrier Sense, Multiple Access with Collision Detect

See **CSMA/CD**.

Channel

The data path between two or more stations, including the communications control capability of the associated stations.

Characteristics

Static information about a component which is kept in either the volatile or permanent database; can include parameters defined for that component with NCP commands.

Circuit

A logical (virtual) communications path providing a communications connection between adjacent nodes. In DECnet, there is one circuit for each point-to-point line, one for each tributary station on a multipoint line, and one for each node on an Ethernet LAN.

Collision

The event when two nodes in a network try to transmit at the same time on the same channel, causing the transmitted data to be unusable.

Collision detect

A signal provided by the Physical layer to the Data link layer to indicate that one or more stations (nodes) are contending with the local station's transmission.

Command node

The node from which an NCP command is issued.

Communications buffers

Buffers allocated at each node for temporary storage of communications data. (In VMS systems, they are also called "executor buffers.")

Communications server

A special-purpose, stand-alone system dedicated to managing communications activities for other computer systems.

Component

An element in the network that can be controlled and monitored. Components include lines, circuits, nodes, logging, modules, and objects.

Configuration database

See **Network configuration database**.



Congestion

A condition occurring when the demands on the network exceed the network's capacity to handle them within a particular time. The condition that arises when arriving traffic exceeds the capability of servers handling the traffic.

Connectivity

The degree to which nodes in the network are interconnected. Full connectivity means all nodes have links to every other node.

Contention control

A scheme of access control used by many networks. Control is distributed among the nodes of the network. Any node that wishes to transmit may do so, accessing the network on a first-come, first-served basis. It does not have to wait its turn. However, it is possible that two nodes may be in contention, or start transmitting at the same time, in which case a collision occurs, and each node must back off and transmit again after waiting a random period of time. An improvement on this scheme is to have the device sense whether the channel is busy before transmitting. Another improvement is to have the device sense while transmitting, to reduce the time needed to detect a collision. The **CSMA/CD** method uses all these techniques.



Control data

Control data is the information used by the network or communicating devices to transfer and process data. Distinguished from user data. For example, control data includes routing information. Control data is considered part of the protocol overhead incurred in communications.

Control station

The network node at the end of a multipoint circuit. The control station controls the tributaries for that circuit.

Controller

See **Line controller**.

Cost

An integer value assigned to a circuit between two adjacent nodes. Packets are routed on paths with the lowest cost. Nodes on either end of a circuit can assign different costs to the same circuit.



Counters

Performance and error statistics kept for a component, such as lines or nodes.

CSMA/CD

Carrier Sense, Multiple Access with Collision Detect: a link management procedure used by the Ethernet. This allows multiple stations to access the broadcast channel at will, avoids contention by means of carrier sense and deference, and resolves contention by means of collision detection and retransmission.

CTERM

Digital Equipment Corporation's Terminal Services Command Terminal protocol. Gives a terminal user the ability to establish a virtual terminal connection to remote DECnet Phase IV systems that also support the protocol.

Data link

A logical connection between two stations on the same circuit on which data integrity is maintained.

Datagram

A unit of data passed between the routing layer and the end communications layer. It is handled independently from all other data. The datagram becomes a packet once the routing header is added.

DECnet

Networking software developed by Digital Equipment Corporation. DECnet is an implementation of **DNA (Digital Network Architecture)**.

Dedicated line

Synonymous with a leased or private line, it is reserved or committed to a specific user or purpose.

Designated router

A routing node on an Ethernet which provides a routing service for the end nodes on the Ethernet.

Destination user

The user to whom the network is delivering information.

Dial-up line

A circuit established by a switched circuit connection.

Distributed processing

A technology that enables the distribution of computing power and storage facilities to user work areas like offices, labs, or desks on factory floors. Distributed processing grew out of the realization that users wasted valuable time by having to wait for answers to their questions from isolated and remotely located mainframe computers. Digital Equipment Corporation is a pioneer in distributed processing.

DNA

Digital Network Architecture: Digital Equipment Corporation's layered data communications protocols. A division of activities into layers required to perform network communications. DNA is compatible with the internationally accepted OSI model. See **OSI**.

Down-line load

The loading of software images onto an unattended node from another node (called the **load host**). For example, Digital's Ethernet communication servers are down-line loaded.

End node

A nonrouting node. It can receive packets addressed to it and send packets to other nodes, but cannot route packets to other nodes.

End user:

1. The data terminal operator at a node in a network who generates or uses information communicated over the network.
2. A computer application program accessing a network to generate or use communicated information.

Ethernet

A local area network that employs coaxial cable as a passive communications medium in a CSMA/CD system to interconnect different types of computers, server products, and office equipment at a local site. No switching logic or central computer is needed to establish or control communications.

Ethernet End node Hello message

This message is used for the initialization and periodic monitoring of end nodes on an Ethernet circuit. Each end node periodically broadcasts an Ethernet End node Hello message to all routers on the Ethernet. The routers use this message to maintain the status (up or down) of end nodes on the Ethernet.

Ethernet Router Hello message

This message is used for initialization and periodic monitoring of routers on an Ethernet circuit. Each router on the Ethernet uses a multicast operation to periodically broadcast an Ethernet Hello message to all other nodes (routers and end nodes) on the same Ethernet. The message contains a list of all routers on the Ethernet from which the sending router has recently received Ethernet Router Hello messages. By exchanging these messages, all routers remain informed of the status of the other routers on the Ethernet.

Event

A network or system-specific occurrence for which the logging component maintains a record.

Event class

A particular classification of events. Generally, this classification follows the DNA architectural layers; some layers may contain more than one class. Class also includes the identification of system-specific events.

Event type

A particular form of event that is unique within an event class.

Executor node

The node at which an NCP command is executed.

Flow control

Hardware or software mechanisms employed in data communications to turn off transmission when the receiving station is unable to store the data it is receiving. In DECnet, the NSP (Network Services Protocol) coordinates the flow of data on a logical link in both directions, from transmit buffers to receive buffers. It ensures that data is not lost, prevents buffer deadlock, and minimizes communications overhead.

Gateway

A special purpose device or software implementation that enables the networks of different vendors to communicate by converting the functions of one vendor's network into functions recognizable by the other vendor's network. A gateway serves the same function as a language translator, who enables persons speaking different languages to communicate.

Gateway server

A communications server that provides access from one type of network to another (networks having different access protocols). For example, Digital Equipment Corporation's SNA Gateway provides access to IBM SNA systems.

Hardware address

For an Ethernet device, the unique Ethernet physical address associated with a particular Ethernet communications controller (usually in read-only memory) by the manufacturer.

Hello and Test message

A message that tests an adjacent node to determine if an adjacency is operational. The routing layer of a node sends this message periodically on non-Ethernet circuits in the absence of other traffic. When a node's routing layer receives this message or any other valid message from an adjacent node, the routing layer starts or restarts a timer. If the timer expires before the routing layer receives another message from the adjacent node, routing considers the adjacent node down.

Hello timer

A circuit parameter that defines the frequency at which routing messages are sent to an adjacent node to keep it aware of the local node's status as a reachable node.

Hertz

A unit of frequency equal to one cycle per second. Cycles are referred to as Hertz (abbreviated Hz).

Hop

The logical distance between two nodes. One hop is the distance from one node to an adjacent node.

Host-based router

Full-function nodes that handle user applications as well as routing. Contrast with **Server-based router**.

Host node

For DECnet, a node that provides services for another node. For example, a load host supplies image files for a down-line load.

Initialization

A start-up procedure between two adjacent nodes wishing to communicate.

Initialization message

This message is used by the routing layer when initializing a non-Ethernet circuit. The message contains information about the node type, required verification, maximum data-link layer receive block size, and routing version.

ISO

International Standards Organization, a standards-defining body based in Switzerland.

Kbps

Kilo (1000) bits per second.

LAN

See **Local area network**.

LAT

Local area transport protocol used for virtual terminal communications on Ethernet LANs. Gives a terminal user the ability to establish a virtual terminal connection to host DECnet Phase IV systems that also support the protocol. The LAT interface is optimized for high terminal I/O performance over an Ethernet, while reducing CPU cycles required by the host system to handle interrupts.

Layer

In networks, layers pertain to the software protocol levels that make up the architecture. Each layer performs certain functions for the layers above it.

Leased line

See **Dedicated line**.

Level 1 router

A node that can send and receive packets, and route packets from one node to another, only within a single area.

Level 2 router

A node than can send and receive packets, and route packets from one node to another, within its own area and between areas. Also known as an area router.

Line

The network management component that provides a distinct physical data path.

Line controllers

Hardware devices (at each node) that manage communications over each line. They handle the physical-link and data-link layers. The speed and efficiency at which these devices handle data can significantly cut short the time required by CPUs for communications functions.

Line speed

The maximum rate at which data can be reliably transmitted over a line. Line speed varies with the capability of the modem or hardware device that performs the transmitting. In general, line speed is described in terms of bits per second (bps).

Link

Any specified relationship between two nodes in a network. A communications path between two nodes.

Load balancing

This is the ability of a system to bring a user job or request to the least loaded resource (of a group of resources providing similar services). For example, when more than one host node connected to a terminal server offers the same service, the server connects the requester of the service to the host with the least load. Thus, load balancing increases access and efficient use of system resources.

Load host

The node from which a system image is down-line loaded to an unattended node.

Local area network (LAN)

A high-speed data communications network that covers a limited geographical area such as an industrial complex or college campus.

Local node

The node where you are located.

Logging

The management facility that collects network events at a logging sink such as a file or console.

Logging console

A logging sink that receives ASCII-encoded events. This is usually a terminal or file.

Logging file

A logging sink which is a file that stores binary-encoded events for later reference.

Logging monitor

A logging sink such as a terminal that receives events as they occur. A system or user program that receives binary-encoded event messages.

Logical link

A temporary connection between processes on source and destination nodes (or between two processes on the same node).

Loop assist test

A loopback test performed on an Ethernet to check the Ethernet circuit between two LAN nodes with the assistance of a third LAN node. This third node, the assistant node, relays test data to or from the destination node, or both ways.

Loop node

Used for loopback testing, a loop node is a local node associated with a specified line.

Maximum cost

A parameter that states a value used by the routing layer to determine when a node is unreachable. If the least costly path to that node exceeds this value, it is unreachable. For correct network operation, this parameter value must not be less than the maximum path cost for the network.

Maximum hops

A parameter that states a value used by the routing layer to decide whether a node is unreachable. If the length of the shortest path to a node exceeds this value, the node is unreachable. For correct network operation, this value must not be less than the network diameter.

Maximum path cost

The value of the path between two nodes of the network that has the greatest routing cost. The routing cost is the least costly path between two nodes.

Maximum path length

The routing distance between two nodes in the network with the greatest routing distance. The routing distance is the length of the least costly path between two nodes.

Maximum visits

The maximum number of nodes through which a packet can be routed before arriving at the destination node. If a packet exceeds the maximum number of visits, it is dropped.

Mbps

Million (mega) bits per second.

Megabits per second

Millions of bits per second.

Message

An ordered collection of data that is intelligible to both the sender and the receiver (at the highest levels of the network).

Modem

“Data sets” at each end of a communications line necessary to convert digital signals transmitted by a computer into analog signals for transmission over a voice-grade line. A modem at the other end converts the analog signal back into digital form for the receiving computer. The term is derived from MODulator/DEMODulator.



Monitor

A software or hardware tool that records events or activities of a system or network and/or analyzes the data. Monitors are used to measure and observe the performance of a network.

Multibuffering

Providing more than one buffer at a time to handle data. For network operations, multibuffering improves performance (if sufficient memory is available). A network can send or process several buffers of data in quick succession. In single-buffering, it handles only one buffer at a time, so subsequent operations requiring buffers must wait until the single buffer is available. Multibuffering is a form of pipelining.


Multicast addressing

An addressing method which sends a message to a group of nodes or all nodes on an Ethernet simultaneously instead of addressing a specific node.

Multidrop line

See **Multipoint line**.

Multiple area network



A network divided into areas, with each area being a group of nodes. Nodes are grouped in areas for hierarchical routing purposes. Hierarchical routing involves an additional level of routing besides the conventional routing. This second level of routing is called "level 2 routing" and is for routing between layers. Routing within an area is called "level 1 routing."

Multipoint line

A single line shared by several nodes. (An Ethernet is a form of a multipoint line, but is more correctly referred to as a multiaccess medium.) Most multipoint configurations include a central node that controls access to the line shared by the other nodes. (Ethernet LANs do not need a central node.)

Network

A group of computers interconnected to share resources and exchange information.

Network configuration database

The combination of both the permanent and the volatile databases. It consists of information about the local node, and all nodes, modules, circuits, lines, and objects in the network.

Network delay

The time it takes to get a unit of data from the source of a transmission to the destination. This usually refers to the delay induced purely by the network and not by system-dependent application processing delays at source and destination nodes.

Network diameter

The distance between the two nodes in the network with the greatest reachability distance. The reachability distance is the length of the shortest path between two nodes.

Network management

The administrative services concerned with managing a network, such as configuring and tuning the network software, monitoring network performance, maintaining network operation, and diagnosing and troubleshooting network problems. In the DNA architecture, the layer containing functions that enable operator control and observation of network parameters and variables.

Node

A point in the network that supports DECnet software and where service is provided, service is used, or communications channels are interconnected.

Node address

The unique numeric identification of a specific node within the network.

Node database

The database that stores information about all remote nodes known to the local system and about the local node itself.

Node name

The alphanumeric string associated with a specific node. It is associated with the node's address, with strict one-to-one mapping.

Nonrouting node

See **End node**.

Operational database

Also called the "volatile database." A memory image containing information about network components. When the system on which the image is running goes down, the database is erased. When the system comes back up, its volatile database inherits the characteristics and parameter values defined in the permanent database.

OSI

Open Systems Interconnection, the only internationally accepted framework of standards for intersystem communication. Developed by the International Standards Organization, OSI is a seven-layer model that covers all aspects of information exchange between two systems.

Out-of-order packet caching

The ability of a node to take packets of a message that have been received out of order and reassemble them into their proper order. This feature must be supported by any destination nodes that receive data that has been split over several paths (see **Path splitting**).

Overhead data

All information other than user information. Overhead data may include: (1) header information on a user message, giving such information as the source and destination address, number of bits included, and so forth; (2) control information for data exchanges between two parties, such as hand-shaking information to initiate and synchronize an exchange, acknowledgment of data received, and so forth; (3) information sent by the network to users for reporting system status or controlling user operations.

Packet

A group of bits, including user data and control data, that are transferred as a unit through a network by the routing mechanism. When stripped of its routing header and passed to the end communications layer, it becomes a datagram.

In packet-switched networks, a packet is the basic unit of transmission.

Packet switching

A data transmission process, using addressed packets that occupy a channel only for the duration of transmission of the packet. Each packet sent by a user can take different routes to the destination. DECnet does not use packet-switching, although it can provide access to packet-switching networks.

Packet switching data network (PSDN)

A set of equipment and interconnecting links that provide a packet-switching communications service to subscribers.

Parameters

Software-related variables that network managers, system managers, and applications programmers can manipulate to optimize performance. In networks, these include the sizes of packets and buffers, the length of certain timers or timeout values, the sizes of certain quotas, and more. Parameters such as these are the “control knobs” that allow managers to tune or adjust the network’s performance.

Path

A possible route a packet takes from a source node to a destination node. The path can comprise a sequence of connected nodes between the source and destination.

Path cost

The sum of the circuit costs along a path between two nodes. A network manager can specify the maximum path cost for a network. In multiple area networks, the manager can specify the maximum cost for a path within an area and for a path between areas.

Path length

The number of hops along a path between two nodes; that is, the number of circuits a packet must travel along to reach its destination.

Path splitting

In DECnet Phase IV, the ability to split the transmission load destined for a single node over several paths of equal total path cost. Any destination node receiving data that has been split over several paths must support out-of-order packet caching.

Permanent database

A file usually stored on a disk containing information about network management components. The database is permanent in that the information remains unchanged after the system or network is brought up after being down. Contrast with **Operational database**.

Physical address

The unique address value that the network associates with a given system on an Ethernet circuit. An Ethernet physical address is defined to be distinct from all other physical addresses on an Ethernet. Initially, a node’s Ethernet physical address is its Ethernet hardware address. After DECnet software is loaded on the node, its physical address becomes the node’s extended DECnet node address. The physical address reverts to the hardware address only when the node is powered off and powered up again. The network software then uses the node’s hardware address until the DECnet software is reloaded on the node.

Piggybacking

A data communications mechanism in which an acknowledgment (of data received) is held until a packet of user data headed for the same destination is transmitted. The acknowledgment is then embedded as part of the protocol header on the packet and “hitches” a free ride. More efficient than sending each protocol message separately, since it significantly reduces overhead.

Pipeline quota

The maximum number of messages or packets that a transmitting node can send without waiting for individual acknowledgments for each successive message.

Pipelining

A technique in which several messages or tasks are handled simultaneously through several stages of a process. This speeds up processing time, as it is much more efficient than processing one message or task at a time through several stages, forcing subsequent messages or tasks to have to wait until the preceding are processed. At the data link layer of a network, pipelining refers to the mechanism by which the transmitting node sends several messages without waiting for individual acknowledgment of each successive message.

Point-to-point

A point-to-point link is a circuit that connects two nodes only. (Contrast with **Multipoint**.) A point-to-point configuration requires a separate physical connection between each pair of nodes.

Polling

A scheme for control of access to a data channel or resource. On a multipoint channel, the control station polls the multipoint circuit’s tributaries to grant the tributaries permission to transmit. The tributaries cannot access the channel on their own. The central node polls each node consecutively, asking if it has anything to transmit. When a node says yes, it is given the right to send.

Protocol

A basic procedure or set of rules that controls the communications between computers. Also, a set of conventions between communicating processes regarding the format and contents of messages to be exchanged. The rules of “etiquette” determine the behavior that the communicating devices can expect of each other.

Protocol overhead

That part of communications data or processing not directly “consumed” by the users, but necessary to successfully bring about the transfer of user information. See **Control data**.

Protocol transparency

The quality in a communications device or system that allows various higher-level protocols to coexist on the same wire. In other words, the protocols are transparent to the device or system.

Queue

A group of data items waiting for service from a device or processor. For example, if incoming data arrives at a rate faster than the device can process it, the data must enter a queue and wait until the device processes the data ahead of it in the queue. Items of data in a queue gain attention from the device according to the priority scheme imposed by the system; for example, first-come, first-served.

Reachable node

A node to which the local node has a usable communications path. The path is usable if it is available and does not exceed the maximum path cost or maximum hops defined for the network.

Reliability

The amount of time a piece of equipment, a system, or a network is available, how often it fails, and how easy it is to maintain. Also, the degree of data integrity maintained in data transmissions.

Remote node

Any node in the network other than the local node.

Repeater

A bidirectional device that amplifies or resynchronizes signals into standard voltages, currents, and timing.

Resource

Any part of the network where traffic is processed, stored, or channeled, such as central processor units (CPUs) at each node, storage disks, communications lines, and line interfaces. An application requires one or more network resources at each stage of execution, and if it does not have access to them, it is delayed or blocked from execution until the necessary resources are available.

Retransmission

A method of error control in which stations receiving messages acknowledge the receipt of correct messages and, on receipt of incorrect messages, either do not acknowledge or acknowledge in the negative. The lack of acknowledgment or receipt of a negative acknowledgment indicates to the sending station that it should transmit the failed message again.

Round trip delay

The time a packet takes to reach the destination and return an acknowledgment (**ACK**) to the source.

Router

A node that can send and receive packets, forwarding them to other nodes.

Router server

A communications server that off-loads routing functions from other nodes in a network. See also **Server-based router**.

Routing

The routing protocol in a network determines the path or **route** along which the data travels to reach its destination. Where multiple paths to a destination are available, routing chooses the best path. See **Adaptive routing**.

Routing configuration parameters

Parameters that specify how the network is to be configured, such as the node type (routing, nonrouting, area routing, etc.), the maximum number of routers, or the maximum node address.

Routing control parameters

Parameters that indirectly control the path that data takes through the network and also control the timing of routing messages. Examples are the circuit cost and routing timer parameters.

Routing database

A database kept on a routing node and which contains regularly updated information about the status of all destination nodes in the area, the status of the communications paths between these nodes and itself, and the paths to other areas in the network. Using the updated information in its routing database, a DECnet routing node can automatically adapt to changing conditions in the network.

Routing node

A router.

Routing queue threshold

A parameter specifying the maximum number of packets that can be queued simultaneously on a circuit before dropping packets.

Routing update message

These messages are exchanged by routing nodes in a network with adaptive routing to keep each router up-to-date about the cost and hops required to reach each node in the network.

Segmentation

The mechanism by which network software divides normal data (handed to the network in buffers) into numbered segments for transmission over a logical link. At the receiving end, the data is reorganized into its original form.

Server

A hardware and software module or device that performs a specific, well-defined service for many users, such as terminal I/O handling (terminal server) or remote file access. A network node that manages the sharing of resources independently of a host processor. You can think of a server as a service.

Server-based router

A routing node dedicated to routing only. Server based routers are designed and optimized for the single function of routing. Routing does not have to contend with other processes or tasks as is the case on host-based routers. Contrast with **Host-based router**.

Sink node

A node used to receive logged events. Logging sinks such as a file or console are located on this node.

SNA (IBM System Network Architecture)

A network for moving IBM mainframe data. Digital Equipment Corporation's SNA Gateway product allows a node not directly connected to an IBM SNA network to access the facilities of the SNA network for terminal access and remote job entry.

Source user

The user who sends information to the network or requests a network process that involves another user (called the **destination user**) who will receive the information or be acted upon by the process.

State

The status of a network component.

Station

A physical termination of a line, having both a hardware and software implementation (a controller and/or a unit). A server, such as a workstation or print station, at which direct user interaction occurs.

Substate

An intermediate circuit state that is displayed in a circuit state display via the network management SHOW or LIST command.

Switched line

A communications link for which the physical path may vary with each use. The dial-up telephone network is an example of a switched line.

Switching

In data communications, identifying and connecting independent transmission links to form a temporary, continuous path from the source to the destination. Contrast with **Dedicated line**.

Synchronous transmission

A mode of transmission in which characters are transmitted at a fixed rate, with the transmitter and receiver synchronized. Greater efficiency is gained by not having to use start and stop bits with each character as in asynchronous transmissions. Synchronous transmissions send a predetermined group of "sync" characters ahead of a long stream of data. The sync characters enable the communicating devices to synchronize with each other in accord with time clocks at each end.

Tap

The entry point where the transceiver connects to an Ethernet cable.

Terminal server

A system that handles terminal operations for host nodes on the LAN. In a LAN, terminal servers can be used to connect terminal users to nodes on the same LAN. Some terminal servers connect users to nodes located off the local LAN. Terminal servers off-load the terminal connection and I/O responsibilities from host nodes on the LAN. Terminal servers reduce the number of direct terminal connections to each host, saving substantial power, packaging, and cabling expenses. Digital Equipment Corporation's terminal servers can connect up to 32 terminals each.

Throughput

In general, throughput is the measure of how much data is sent, or can be sent, between two points in a specified unit of time. Throughput is often used in either of two contexts: (1) rated throughput, which refers to the bandwidth or capacity of a component, or (2) real throughput, which refers to actual measured throughput.

Throughput, gross

Measures the total amount of data transferred in time, including the protocol overhead as well as the "consumable" user part of the data. Gross throughput is also called "total throughput." Contrast with **Throughput, net**.

Throughput, net

Measures the amount of user data that is transferred, not counting network or protocol data. By user data is meant that which the user directly uses or "consumes." Network or protocol control data is useful, but not directly consumed by the user. Users are interested only in sending or receiving the user data. (Also called "net throughput," "effective throughput," and "productive throughput.") Contrast with **Throughput, gross**.

Throughput, network

The sum of the lengths of all terminating messages received on all nodes across a network per second, where a terminating message is any packet that has reached its destination node. Also called the "aggregate throughput."

Throughput, rated

The theoretical or potential throughput of a component, based on its capacity, regardless of actual, applied conditions that will detract from it in real life. Contrast with **Throughput, real**.

Throughput, real

Describes the measured processing rate or transfer rate of the user information. This meaning of throughput is sometimes called the "transfer rate of information blocks (TRIB)." It is less than capacity (that is, less than the rated throughput), because of the time required by the application or network to prepare the data for transmission. Intervals between data during the measured interval of time detract from the maximum, possible throughput.

Topology

The physical shape or arrangement (geometrical structure) of interconnected nodes in a network. Where possible, the topology should be based on the nature of the traffic flow in the network; otherwise, performance can be hampered significantly.

Traffic

The measurement of data flow, volume and velocity over a communications link.

Transceiver

A device required in baseband networks that takes the digital signal from a computer or terminal and imposes it on the baseband medium.

Transfer rate, data

The actual, measured rate at which data is transferred, such as bits per second. This does not take into account intervals between transmissions of bits. Therefore, the measured transfer rate is less than the maximum transfer rate. See **Throughput, real**.

Transfer rate, maximum

The signaling capacity of a component, or its theoretical throughput, disregarding any of the limiting factors that come into play in real life (real applications). It is analogous to the EPA rating for automobiles. Synonymous with **line speed** or **signaling rate** of a communications line. Contrast with **Transfer rate, data**. See **Throughput, rated**.

Transparency

The quality in a network that enables users to access and transfer information without having to know how the network operates. For example, in DECnet the routing function is transparent to users.

Transparency, protocol

See **Protocol transparency**

Tributary station

In DECnet multipoint links, a station on a multipoint line that is not a control station.

Up-line dump

A function that allows an unattended node to dump its memory to a file on another node.

User

A user can be a human operator at a terminal, an unattended device, or an application program. Network communications usually consist of two operators at interactive terminals, two programs swapping messages, or an operator who wants to run a remote program or open a remote file.

User information

The information (message) that the source user wishes to convey to the destination user.

Utilization

A measure that describes the percentage of time that a resource is busy. It indicates how much use a resource is getting.

Verification message

This message is sent along with the Initialization message on a non-Ethernet circuit if verification is required.

Virtual terminal

A terminal physically connected to one node in a network but, by using network software, can access other nodes and perform the same activities on them as a directly connected terminal. (On VAX/VMS systems, this is done by the SET HOST command.)

Volatile database

See **Operational database**.

WAN

See **Wide area network**.

Wide area network

A data network that covers a wide geographical area, used for long-distance communications. Nodes can be separated by thousands of miles.

Window

A range of data packets authorized to be in transit across a network.

Window flow control

A form of flow control based on the window size. When the window size is reached, further transmission stops until the number of outstanding packets drops below the window. The optimum window size brings maximum throughput. If the window size is too small, throughput is restricted. If the window size is too large, throughput will start to decrease when the network becomes congested.

Window size

The maximum number of data packets permitted in transit at one time.

X.25 Gateway Access Protocol

In DECnet, a protocol to allow a node not directly connected to a public data network to access facilities of that network through an intermediary gateway node. X.25 is the protocol standard for communication in packet-switched data networks.

Index

A

Access
 see Network access
Access control verification, 2-7
 rejected, 5-5
Adaptive routing, 2-6
 DECnet, 2-1 to 2-2
 defined, 1-2
Adjacency, 2-8
Adjacency rejected (event message),
 5-3, 5-4
Adjacent node
 defined, 1-1
Adjacent node block size too small
 (event message), 5-6
Applications
 file transfer, 4-4, 4-6
 local, 4-2
 network, 3-1, 4-2
 remote terminal, 4-4, 4-6
Area, 2-4
Area leakage, 5-10 to 5-12
Area number, 2-6, 4-3
Area router
 see Level 2 router
Area routing, 2-13
 see also Multiple area network
 and DECnet Phase III, 5-10
 and Ethernet, 5-12
 benefits of, 2-6, 3-11

Area routing (cont.)

 configuration guidelines for, 4-13
 defined, 2-4
 parameters, 2-6, 5-10
Area routing problems, 5-5, 5-10 to 5-15
Asynchronous line, 3-5, 3-7
Asynchronous router, 3-8, 3-9

B

Bandwidth efficiency, 3-16
Bridge, 3-4, 3-8, 3-12, 3-14 to 3-15, 4-10
 and maximum broadcast nonrouters
 parameter, 4-10
 and maximum broadcast routers
 parameter, 4-12
 and router update messages, 4-16
 compared with router server, 3-16
Broadcast address, 3-6
Broadcast nonrouter, 4-9, 4-10, 5-3, 5-4
 see also Maximum broadcast nonrouter
 parameters, Ethernet end nodes
Broadcast router, 4-10, 4-13
 see also Maximum broadcast router
 parameters, Ethernet routing nodes
Broadcast routing timer, 2-9
Buffer allocation problems, 5-6, 5-7
Buffer size parameter, 4-4 to 4-5, 5-6
 minimum acceptable size, 5-6
Buffers, 2-7
Bus topology, 3-4
Bytes sent (counter), 5-8

C

- Carrier Sense, Multiple Access with Collision Detect
 - see* CSMA/CD
- Centralized processing, 4-8
- Circuit, 1-1
- Circuit bandwidth, 4-15
- Circuit characteristics, 3-2
- Circuit congestion, 5-8
- Circuit cost, 2-4, 2-10 to 2-11
 - controlling data flow with, 2-4, 2-10
 - defining, 2-4
 - formula for determining, 4-5
- Circuit down (event message), 5-3, 5-4, 5-6
- Collision, 3-6
- Communications problems
 - see* Routing problems
- Communications server, 3-12
- Computer interconnect (CI), 3-5
- Congestion, 4-15, 5-7
 - control of, 2-4
 - prevention of, 2-13, 4-8
 - troubleshooting, 5-8
- Connectivity and routing, 1-2
- Connectivity problems, 5-1
- Control station, 3-5
- Cost, 1-2, 2-4
 - see also* Circuit cost, Path cost
- Counters, 2-7, 5-9
- CPU overload, 5-9
- CSMA/CD, 3-6

D

- Data link problems, 5-5
- Data link technologies, 3-5
- DDCMP, 2-6, 3-5, 3-7 to 3-10
- DECnet, 3-1, 4-1
 - data link technologies, 3-5 to 3-10
 - environments, 3-3 to 3-4
- DECnet Phase III, 3-11
 - and area routing, 5-10
 - and maximum node address, 4-4, 5-2
 - coexistence with Phase IV, 5-14
 - in multiple area networks, 4-14, 5-10
 - node characteristics, 3-11
 - routing messages, 2-9

- DECnet Phase IV, 2-4
 - coexistence with Phase III, 5-14
 - node characteristics, 3-11
 - routing enhancements, 2-13 to 2-15
 - routing messages, 2-9
- DECnet routing, 1-2, 2-1 to 2-15
- DECnet routing node, 1-1
- DECnet software version, 5-4
- Designated router, 4-13

- defined, 3-6
- performance problems with, 5-9
- Designated Router Hello message, 2-8
- Dial-up line, 3-3
- Digital Data Communications Message Protocol
 - see* DDCMP
- Digital Network Architecture
 - see* DNA
- Distributed network, 3-1
- Distributed processing, 4-8
- DNA
 - layers, 3-12 to 3-13
 - routing, 1-1

E

- End node, 3-12, 4-15
 - and circuit cost, 2-10
 - benefits of, 1-6
 - defined, 1-1
 - Ethernet, 4-9 to 4-10
- Ethernet, 1-3, 2-6
 - addressing, 3-6
 - area routing problems on, 5-12
 - baseband, 3-4
 - broadband, 3-4
 - data link technology, 3-6 to 3-7
 - designated router on, 3-6
 - end nodes, 3-6, 4-9, 4-9 to 4-10, 5-3, 5-4
 - multiple area network on, 3-4, 4-9, 5-12
 - network access to, 3-5
 - routing, 3-4, 3-6
 - routing nodes, 3-6, 3-8, 4-9, 4-10 to 4-13, 4-15
 - performance problems with, 5-9
- Ethernet End Node Hello message, 2-8
- Ethernet LAN, 3-4, 4-6

Ethernet LAN (cont.)
and path costs, 4-6
and routing nodes, 4-15
configuration guidelines for, 4-9 to 4-13
data link technology, 3-5
distance limitations, 3-14
extension of, 3-12 to 3-18
Ethernet Router Hello message, 2-8
Event logging, 4-2, 5-9
Event messages, 2-7
Extended LAN, 3-12 to 3-18

F

Full-function routing node
see Host-based router

G

Gateway, 3-4, 3-15

H

Hello and Test message, 2-8
Hop, 1-2, 2-11, 2-12
defined, 2-2
Host-based router
defined, 1-5
performance problems with, 5-7

I

IBM SNA network, 3-4
Initialization message, 2-8
International Standards Organization
see ISO
ISO
layers, 3-12 to 3-13
routing, 1-1

L

LAN, 1-3, 3-3, 3-3 to 3-4
configuration guidelines for, 4-2
defined, 3-1
LAT protocol, 3-15
Leased line, 3-3
Level 1 router, 2-4, 2-6
routing problems with, 5-5
Level 1 routing update message, 2-13
Level 2 router, 2-4, 2-6, 5-12

Level 2 routing, 4-8, 4-14
Level 2 routing update message, 2-13

Line

defined, 1-1
dial-up, 3-3
leased, 3-3
Line characteristics, 3-2
Line congestion, 5-8
Line speed, 3-3
mismatch, 5-5
Local area network

see LAN

Local Area Transport protocol
see LAT protocol

Local buffer errors (counter), 5-7

Lost packet problems, 5-6

M

Maximum area parameter, 2-6, 4-4, 5-10
Maximum broadcast nonrouters parameter,
4-9 to 4-10
too large, 5-3
too small, 5-3
Maximum broadcast routers parameter,
4-10 to 4-13
too small, 5-4
Maximum node address parameter, 4-3 to 4-4
formula for determining, 4-4
too small, 5-2
Maximum path cost parameter, 2-12
between areas, 2-6, 2-12, 4-8, 5-10
defining, 4-7
too small, 5-4
Maximum path length parameter, 2-3, 2-12
between areas, 2-6, 2-12, 4-8, 5-10
defining, 4-7
too small, 5-4
Maximum routers parameter, 4-13
Maximum visits parameter, 2-7, 2-12,
2-13, 4-7
Memory problems, 5-7
Modem, 3-3
Monitoring tools, 4-2
Multicast address, 3-6
Multiple area network, 2-4, 2-7, 3-1, 3-10
to 3-11
see also Area routing

Multiple area network (cont.)
 and maximum broadcast nonrouters
 parameter, 4-10
 and maximum broadcast routers
 parameter, 4-11
 area leakage, 5-10
 configuration guidelines for, 4-13,
 A-1
 converting standard network into,
 A-1
 design of, A-4
 Ethernet, 3-4, 4-9, 5-12
 maximum area parameter for, 4-4
 node address specification for, 4-3
 parameters, 2-12
 routing, 2-6, 2-13
 routing nodes, 4-16
 routing parameters, 2-6
 routing problems with, 5-10
Multipoint line, 1-3, 2-6, 3-5

N

NCP, 4-2

Network

 centralized, 3-1, 3-2
 distributed, 3-1, 3-2
 point-to-point, 1-2
 technologies and topologies, 3-1 to
 3-18

Network access, 3-2, 3-5

 Ethernet, 3-5, 3-6
 multipoint, 3-5
 to non-DECnet networks, 3-4
 WAN to Ethernet, 3-4

Network applications, 3-1

Network buffer size

see Buffer size

Network configuration, 3-1 to 3-18

 factors affecting, 3-1
 for optimal routing performance,
 4-1 to 4-16
 general guidelines for, 4-2 to 4-9

Network configuration database, 4-1

Network connectivity, 1-2

Network Control Program

see NCP

Network delay

 and path length, 4-6

Network delay (cont.)

 excessive, 5-10

Network Management Control

 Center/DECnet Monitor

see NMCC/DECnet Monitor

Network manager, 1-1

 and configuring the network, 2-4, 3-1,
 3-3, 3-10, 4-1
 and control of data flow, 2-4, 2-10, 4-8
 and troubleshooting, 2-7

Network monitoring tools, 4-2

Network performance, 2-10, 3-2, 3-3, 4-15

see also Routing performance

Network problems, 5-1 to 5-15

 and connectivity, 5-1

Network reliability, 1-4, 3-2, 3-17

Network traffic, distribution of, 4-8

NMCC/DECnet Monitor, 4-2

Node address, 4-3

 maximum, 4-3, 5-2

Node address duplication, 5-1

Node characteristics, 3-2, 3-11 to 3-12, 4-1

Node identification, specifying, 4-3

Node names, 4-3

Node numbers, specification of, 4-3

Node out of range packet loss (event
 message), 5-2

Node unreachable packet loss (event
 message), 5-4

Nonadjacent node, 1-1

Nonrouting node

see End node

O

Out-of-order packet caching, 2-14, 5-10

Oversized packet loss (event message), 5-6

P

Packet switching data network (PSDN),
 3-3, 3-4

Partial routing update loss (event message),
 5-3

Path control parameters, 2-12

Path cost, 2-12

 between areas, 2-6, 2-12, 4-8

 defined, 2-4

 maximum, 4-7, 5-4

Path length, 4-6 to 4-8

Path length (cont.)
 and network delay, 5-10
 between areas, 2-6, 2-12, 4-8
 defined, 2-2
 maximum, 2-3, 2-12, 4-7, 5-4
Path splitting, 2-4, 2-13 to 2-15, 4-8, 5-10
Path, choice of, 1-2, 2-1, 2-4
Path-splitting parameter, 2-13
Performance
 see also Network performance
 bridges, 3-13
 of server-based routers, 3-13
Performance problems, 5-6 to 5-10
 in end-to-end communications, 5-7
 system, 5-7
Personal computers, 3-9
Point-to-point line, 2-6, 3-5, 3-7
Point-to-point network, 1-2 to 1-5
Polling, 3-5
Postal telegraphic and telephone authority (PTT), 3-3
Problems
 see Area routing problems, Data link problems, Lost packet problems, Memory problems, Performance problems, Routing
Protocol, 3-16
Protocol transparency, 3-15, 3-16
PSDN
 see Packet switching data network
PTT
 see Postal telegraphic and telephone authority

R

Rated throughput, 5-9
RBMS
 see Remote Bridge Management Software
Reachable areas, 5-10
 defined, 2-7
Reachable nodes, 2-12
 defined, 2-7
Receive password, 5-5, 5-12
Reliability, 1-4
Remote Bridge Management Software (RBMS), 4-16
Remote buffer errors (counter), 5-7

Repeater, 3-4, 3-12
 defined, 3-14
Response time, 3-1, 3-2, 4-6
 see also Network delay
Router
 see Asynchronous router, Routing node, Synchronous router
Router server
 see Server-based router
Routing, 1-2
 see also Adaptive routing, Area routing and connectivity, 1-2, 1-3 and user applications, 1-5, 1-6
 benefits of, 1-2 to 1-5
 DECnet, 2-1 to 2-15
 defined, 1-1
 Ethernet, 3-6
 features, 1-2, 2-6 to 2-7
 introduction to, 1-1 to 1-7
 on full-function nodes, 1-5
 on host-based routers, 1-5
 on server-based routers, 1-5, 1-6
 overhead, 3-11, 3-17
Routing configuration parameters, 2-10
Routing control message, 2-8 to 2-9
Routing control parameters, 2-9
Routing database, 1-1, 1-2, 2-2, 2-7, 4-3, 4-6, 4-10, 4-13
Routing layer, 1-1, 2-4
Routing node, 3-12
 see also Host-based router, Server-based router
 benefits of, 1-2 to 1-5
 CPU overload on, 5-9
 defined, 1-1
 Ethernet, 4-10 to 4-13
 full function
 see Host-based router
 host-based, 1-5, 5-7
 performance problems with, 5-6
 server-based, 1-5
 when and where to include, 4-15 to 4-16
Routing overhead, 4-9, 4-15
Routing parameters, 2-9 to 2-13, 5-1
Routing path, 2-2
Routing performance optimization, 4-1 to 4-16
Routing performance problems, 5-8

- Routing priority parameter, 3-7
- Routing problems, 5-1 to 5-15
 - connectivity, 5-1
 - general, 5-1 to 5-10
- Routing queue threshold parameter, 5-7
- Routing timer, 2-9
- Routing traffic, calculating amount of, 5-9
- Routing update message, 2-2, 2-7, 2-8
 - DECnet Phase III, 2-9
 - DECnet Phase IV, 2-9
 - level 1, 2-13
 - level 2, 2-13
 - multicast, 2-13

S

- Satellite, 3-17
- Seconds since last zeroed (counter), 5-8
- Segment buffer size parameter, 4-4
- Segmented routing messages, 2-9
- Server-based router
 - compared with bridge, 3-15 to 3-18
 - defined, 1-5
- SNA
 - see* IBM SNA network
- Synchronous line, 3-5, 3-7
- Synchronous router, 3-8
- System manager, 4-1
- System performance degradation, 5-7

T

- Throughput, 3-1, 3-5, 4-5, 5-9
- Throughput requirements, 4-9
- Traffic
 - see* Routing traffic
- Transit congestion loss (counter), 5-8
- Transit packets received (counter), 5-9
- Transit packets sent (counter), 5-8
- Transmit congestion loss (counter), 5-7
- Transmit password, 5-5, 5-11
- Tributary, 3-5

U

- User buffers unavailable (counter), 5-7
- Utilization, 2-13

V

- VAX ETHERnim, 4-2
- Verification message, 2-8
- Verification password, 2-7
 - see also* Access control verification,
 - Receive password, Transmit password
- Verification reject (event message), 5-5
- Version skew (event message), 5-4
- Virtual circuit, 3-7

W

- WAN, 1-3, 3-3
 - configuration guidelines for, 4-2
 - data link technologies for, 3-5
 - defined, 3-1
- Wide area network
 - see* WAN

X

- X.25, 2-6, 3-4, 3-5, 4-5

HOW TO ORDER ADDITIONAL DOCUMENTATION

DIRECT TELEPHONE ORDERS

In Continental USA
and Puerto Rico
call 800-258-1710

In Canada
call 800-267-6146

In New Hampshire
Alaska or Hawaii
call 603-884-6660

ELECTRONIC ORDERS (U.S. ONLY)

Dial 800-DEC-DEMO with any VT100 or VT200
compatible terminal and a 1200 baud modem.
If you need assistance, call 800-DEC-INFO.

DIRECT MAIL ORDERS (U.S. and Puerto Rico*)

DIGITAL EQUIPMENT CORPORATION
P.O. Box CS2008
Nashua, New Hampshire 03061

DIRECT MAIL ORDERS (Canada)

DIGITAL EQUIPMENT OF CANADA LTD.
940 Belfast Road
Ottawa, Ontario, Canada K1G 4C2
Attn: A&SG Business Manager

INTERNATIONAL

DIGITAL
EQUIPMENT CORPORATION
A&SG Business Manager
c/o Digital's local subsidiary
or approved distributor

Internal orders should be placed through the Software Distribution Center (SDC),
Digital Equipment Corporation, Northboro, Massachusetts 01532

*Any prepaid order from Puerto Rico must be placed
with the Local Digital Subsidiary:
809-754-7575



READER'S COMMENTS

What do you think of this manual? Your comments and suggestions will help us to improve the quality and usefulness of our publications.

Please rate this manual:

	Poor			Excellent	
Accuracy	1	2	3	4	5
Readability	1	2	3	4	5
Examples	1	2	3	4	5
Organization	1	2	3	4	5
Completeness	1	2	3	4	5


Did you find errors in this manual? If so, please specify the error(s) and page number(s).

General comments:

Suggestions for improvement:

Name _____ Date _____
Title _____ Department _____
Company _____ Street _____
City _____ State/Country _____ Zip Code _____

DO NOT CUT FOLD HERE AND TAPE



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY LABEL
FIRST CLASS PERMIT NO. 33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

digital

**Networks and
Communications Publications**
550 King Street
Littleton, MA 01460-1289



DO NOT CUT FOLD HERE



CUT ALONG DOTTED LINE