
NETWORK TROUBLESHOOTING GUIDE

Order Number: EK-339AB-GD-002

This manual provides an overview of network troubleshooting tools and methodologies, and detailed troubleshooting procedures for specific network problems for DECnet Phase IV and TCP/IP networks. It addresses networks consisting of computers running the VMS and ULTRIX operating systems.

Revision/Update Information: This is a revised manual.

First Printing, July 1989
Second Printing, August 1990

© Digital Equipment Corporation 1990.

The Reader's Comments form at the end of the hardcopy version of this document requests your critical evaluation for use in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

DDCMP, DEC, DECnet, DECnet-DOS, DECSA, DECconnect, DECrouter, DECserver, DECstation, DELNI, DELUA, DEQNA, DEUNA, DNA, FDDI, LAT, METROWAVE, MicroVAX I, MicroVAX II, MicroVAX, RAINBOW, RSX, RdB/VMS, ReGIS, ThinWire, VAX 6000, VAX 8200, VAX 8250, VAX 8300, VAX 8350, VAX 8500, VAX 8530, VAX 8550, VAX 8600, VAX 9000, VAX ETHERnim, VAX Notes, VAX VTX, VAX, VAX-11/750, VAX-11/780, VAXcluster, VAXstation, VMS, VT, and the DIGITAL logo.

The following are third-party trademarks:

AppleTalk is a trademark of Apple Computer, Inc. IBM is a registered trademark of International Business Machines Corporation. Telenet is a trademark of GTE Telenet Communication Corporation. TYMNET is a registered trademark of TYMNET, Inc.

This document is available both in printed and Bookreader format.

This document was prepared using VAX DOCUMENT Version 2.1.

Contents

Preface	xiii
1 Overview of Data Communications Networks	
What Is a Data Communications Network?	1-1
What Are the Components?	1-1
What Does a Data Communications Network Look Like?	1-2
Star Topology	1-3
Ring Topology	1-3
Bus Topology	1-4
Mesh Topology	1-4
Tree Topology	1-5
Local Area Network Topology	1-5
Wide Area Network Topology	1-6
What Is a DECnet Network?	1-6
What Does a DECnet Network Look Like?	1-7
What Communications Media Does DECnet Use?	1-7
What is Ethernet/802.3?	1-8
How Do Systems Communicate over DECnet?	1-8
How Does DECnet Route Messages?	1-9
How Is DECnet Software Structured?	1-10
What Systems Can Communicate over the Network?	1-12
What Network Environments Does DECnet Support?	1-12
Local Area Networks	1-13
Wide Area Networks	1-13
What Is an Internet Network?	1-17
What Does an Internet Look Like?	1-18
What Is TCP/IP and How Is It Structured?	1-19
What Are Addressing, Naming, and Routing?	1-20
Addressing	1-20
Common Internet Address Notation	1-21
Network Classes	1-22
Subnet Addressing	1-24
Advantages of Subnet Addressing	1-27
Subnets and Internet Addresses	1-27
Address Masks	1-28
Naming	1-29
Domain Name System	1-29
Name Servers	1-31
Routing	1-32
IP Routing Tables	1-32
Direct Routing	1-33
Indirect Routing	1-33
Routing In the Presence of Subnets	1-34

Network Protocols	1-36
Internet Protocol	1-36
Routing Protocols	1-37
Routing Information Protocol	1-37
Internet Control Message Protocol (ICMP)	1-37
Transport Protocols	1-38
Transmission Control Protocol	1-38
User Datagram Protocol	1-39
Other Protocols	1-39
Address Resolution Protocol	1-39
Simple Network Management Protocol	1-40

2 Network Troubleshooting Methodology

Knowing Your Network	2-1
Topology	2-1
Architecture	2-2
Using DECnet Architectural Information in Problem Solving	2-2
Using TCP/IP Architectural Information in Problem Solving	2-3
Performance	2-4
Typical Use	2-4
Overview of the Network Troubleshooting Methodology	2-5
Applying The Methodology to Network Problems	2-6
Detecting a Problem	2-9
Obtaining and Refining the Problem Statement	2-9
Gathering Information	2-9
Analyzing, Interpreting, and Classifying Information	2-11
Extent of the Problem	2-12
Node or Host Problems	2-12
LAN Problems	2-12
WAN Problems	2-13
Types of Errors	2-13
Hard Errors	2-13
Inconsistent Errors	2-13
Intermittent Errors	2-14
Transient Errors	2-14
Sources of Errors	2-14
User Errors	2-14
Hardware Errors	2-14
Software Errors	2-15
Configuration Errors	2-15
Interoperability Errors	2-15
Isolating the Source of the Problem	2-15
Isolating to the Node or Host Level	2-16
Isolating to the LAN Level	2-18
Isolating to the WAN Level	2-18
Solving the Problem	2-19
Verifying the Solution	2-20
Cleaning Up	2-20
Verifying the Solution Again	2-20
Documenting the Problem and Solution	2-20

3 Network Management and Troubleshooting Tools

arp command	3-2
Overview	3-2
Required Resources & Privileges	3-3
How to Use	3-3
Recommended Use	3-3
Authorize Utility	3-3
Overview	3-3
Required Resources & Privileges	3-3
How to Use	3-3
Recommended Use	3-4
Breakout Boxes	3-5
Overview	3-5
Required Resources & Privileges	3-5
How to Use	3-5
Recommended Use	3-5
DEC Extended LAN Management Software	3-5
Overview	3-5
Required Resources & Privileges	3-6
How to Use	3-6
Recommended Use	3-6
DECmcc Management Station for ULTRIX	3-7
Overview	3-7
Required Resources & Privileges	3-8
How to Use	3-8
Recommended Use	3-8
LAN Traffic Monitor	3-8
Overview	3-8
Required Resources & Privileges	3-9
How to Use	3-9
Recommended Use	3-9
LAT Control Program	3-10
Overview	3-10
Required Resources & Privileges	3-10
How to Use	3-11
Recommended Use	3-11
netstat command	3-12
Overview	3-12
Required Resources & Privileges	3-12
How to Use	3-12
Recommended Use	3-13
Network Control Program	3-13
Overview	3-13
Required Resources & Privileges	3-15
How to Use	3-15
Recommended Use	3-16
NMCC/DECnet Monitor	3-18
Overview	3-18
Required Resources & Privileges	3-20
How to Use	3-20
Recommended Use	3-20
NMCC/VAX ETHERnim	3-20
Overview	3-20
Required Resources & Privileges	3-21

How to Use	3-21
Recommended Use	3-22
ping command	3-22
Overview	3-22
Required Resources & Privileges	3-22
How to Use	3-22
Recommended Use	3-23
Protocol Analyzers	3-23
Overview	3-23
Required Resources & Privileges	3-23
How to Use	3-23
Recommended Use	3-23
syslog daemon	3-24
Overview	3-24
Required Resources & Privileges	3-24
How to Use	3-24
Recommended Use	3-25
Terminal Server Manager Software	3-26
Overview	3-26
Required Resources & Privileges	3-26
How to Use	3-26
Recommended Use	3-27
traceroute command	3-27
Overview	3-27
Required Resources & Privileges	3-28
How to Use	3-28
Recommended Use	3-30
uerf command	3-30
Overview	3-30
Required Resources & Privileges	3-30
How to Use	3-30
Recommended Use	3-32

4 Resources for Troubleshooting

Log Files	4-1
VMS OPCOM and the Operator Log File	4-1
VMS Netserver Log File	4-2
VMS Error Log File	4-2
VMS Accounting Log File	4-3
ULTRIX Error Log Files	4-3
Routing Path Trace Procedures	4-3
NCP Routing Path Trace Procedure	4-4
netstat Routing Path Trace Procedure	4-5
NCP Loopback Tests	4-6
NCP Loopback Test Results	4-8
Types of Loopback Tests	4-8
Node-Level Tests	4-9
Remote Loopback Test	4-9
Local and Remote Loopback Tests Using a Loop Node Name	4-10
Local-to-Remote Testing	4-11
Local Loopback Test	4-12

Circuit-Level Tests	4-13
Software Loopback Test	4-14
Controller Loopback Test	4-14
Modem Loopback Tests	4-15
Using Loopback Tests to Check Circuitry	4-15
Reachability Tests for TCP/IP Networks	4-17
NCP Counters	4-18
Errors Applicable to Node Problems	4-19
Errors Applicable to LAN Problems	4-19
Errors Applicable to WAN Problems	4-19
Errors Applicable to Node, LAN, and WAN Problems	4-20
Formulas for Understanding Counters	4-20
Packet Rate	4-20
Circuit Quality	4-21
Transit Congestion Loss	4-22
Ethernet Line Statistics	4-22
Retransmissions	4-23
Routing Overhead	4-23
ULTRIX Counters	4-23
DECnet Events	4-23

5 Network Troubleshooting Procedures

Organization	5-1
Troubleshooting Notes	5-3
Aborted Service Request	5-4
Adjacency Rejected/Adjacency Up	5-8
Asynchronous DECnet Problems	5-10
Babbling Device	5-14
Broadcast Storm	5-16
Circuit State Problems	5-20
Connect Failed, Access Control Rejected	5-23
Connect Failed, Unrecognized Object	5-29
Connection Timed Out	5-31
Device Not Mounted	5-36
Dialup Problems	5-38
Host Is Unreachable	5-42
Insufficient Resources at Remote Node	5-47
Invalid Parameter Value	5-51
LAN Bridge Cannot Downline Load	5-53
LAN Segment Communication Problem	5-58
LAT Port Hung	5-62
LAT Print Queue Problems	5-65
Line Synchronization Lost	5-74
Login Incorrect	5-78
Login Information Invalid	5-80
Network Is Unreachable	5-87
Network Object Unknown	5-91
Network Partner Exited	5-94
Node Out of Range Packet Loss	5-98

Partial Routing Update Loss	5-101
Partitioned Area	5-102
Permission Denied	5-106
Remote Node Is Not Currently Reachable	5-109
Terminal Server Connection Failures	5-114
Unknown Host	5-119
Verification Reject	5-123

A Ethernet Configuration Guidelines

Baseband Ethernet	A-1
ThinWire Ethernet	A-1
Fiber-Optic Cable	A-2
DELNI Local Network Interconnect	A-2
DEMPR ThinWire Multiport Repeaters	A-2
DEREP Ethernet Repeaters	A-2
Bridges	A-3

B RFC Request Process

Obtaining RFCs Through The Internet File Transfer Protocol	B-1
Logging In To The Remote Host	B-1
Using Public Directories on The Remote Host	B-2
Obtaining an Index of Available RFCs	B-2
Copying RFC Files	B-3
Logging Out of The Remote Host	B-3
Obtaining RFCs Through The NIC Automatic Mail Service	B-3
Obtaining RFCs by Telephone Request	B-3

Glossary

Index

Examples

3-1	Authorize Utility Example	3-4
3-2	LAN Traffic Monitor Example	3-10
3-3	VMS LATCP Example	3-11
3-4	ULTRIX lcp Example	3-12
3-5	netstat -i Example	3-13
3-6	VMS Network Control Program Example	3-16
3-7	ULTRIX Network Control Program Example	3-18
3-8	Long Output from the ping Command	3-23
3-9	syslog Example	3-25
3-10	Terminal Server Manager Example	3-27
3-11	uerf Example	3-32
4-1	netstat Routing Path Trace	4-6
5-1	SYSGEN SHOW/DEVICE Display	5-11

Figures

1-1	Point-to-Point Connections	1-2
1-2	Ethernet Connections	1-2
1-3	Star Topology	1-3
1-4	Ring Topology	1-4
1-5	Bus Topology	1-4
1-6	Mesh Topology	1-5
1-7	Tree Topology	1-5
1-8	Network Nodes, Circuits, and Lines	1-9
1-9	DECnet Network Architecture (DNA) Layers and Protocols	1-11
1-10	Small Local Area Network Configuration	1-13
1-11	Large Local Area Network Configuration	1-14
1-12	DDCMP Connections	1-15
1-13	Wide Area Network Connections	1-16
1-14	Large Integrated DECnet Configuration	1-17
1-15	Two Networks Interconnected to Form an Internet	1-19
1-16	Multiple Networks Interconnected to Form an Internet	1-19
1-17	Internet Protocol Model	1-19
1-18	Parts of an Internet Address	1-21
1-19	Internet Address Fields	1-21
1-20	Common Internet Address Notation	1-21
1-21	Internet Address Format for Network Classes	1-23
1-22	Internet Network Hierarchies	1-25
1-23	Internet Without Subnets	1-26
1-24	Internet with Subnets	1-27
1-25	Internet Subnet Address Parts	1-28
1-26	Internet Name Domains	1-31
1-27	Name Server Process	1-32
1-28	Internet Routing Process	1-34
1-29	Subnet Routing Examples	1-36
2-1	Network Troubleshooting Methodology	2-6
2-2	Unreachable DECnet Node or TCP/IP Host	2-7
2-3	Isolating the Source of the Problem	2-17
4-1	NCP Routing Path Trace Procedure	4-5
4-2	Components and Points of Loopback Testing	4-7
4-3	Remote Loopback Test	4-10
4-4	Local-to-Remote Loopback Test Using a Loop Node Name	4-11
4-5	Local-to-Local Loopback Test Using a Loop Node Name	4-12
4-6	Local Loopback Test	4-13
4-7	Software Loopback Test	4-14
4-8	Controller Loopback Test	4-15
4-9	Modem Loopback Tests	4-17
5-1	ULTRIX Access Control	5-24
5-2	VMS Access Control	5-81
5-3	Area Partitioned Due to Multiple Failures	5-103
5-4	Area Partitioned Due to Configuration Weaknesses	5-104

Tables

1-1	DECnet Layers	1-11
1-2	TCP/IP Layers	1-20
1-3	Sizes of Internet Network Classes	1-23
1-4	Address Ranges for Network Classes	1-24
1-5	Internet Address Fields Available For Subnets	1-29
1-6	Internet Domains	1-30
2-1	Tools for Node or Host Problems	2-17
2-2	Tools for LAN Problems	2-18
2-3	Tools for WAN Problems	2-19
3-1	Network Management and Troubleshooting Tools	3-1
3-2	ULTRIX netstat Command Options	3-13
3-3	Options for the ping Command	3-22
3-4	ULTRIX traceroute Command Options	3-29
3-5	ULTRIX uerf Command Options	3-30
5-1	Network Problems and Extent of Disturbance on the Network	5-2
5-2	NCP Executor Proxy Access Parameters	5-25
5-3	LAN Bridge 100 or 150 Indicator Lights	5-54
5-4	LAN Bridge 100 or 150 Switch Settings	5-57
5-5	Print Queue States and Conditions	5-67
5-6	NCP Proxy Access Parameters	5-82
B-1	Public Directories on NIC.DDN.MIL	B-2

Preface

The *NETWORK TROUBLESHOOTING GUIDE* provides an overview of network troubleshooting tools and methodologies, and network troubleshooting procedures for specific network problems.

This version of the *NETWORK TROUBLESHOOTING GUIDE* discusses DECnet Phase IV networks and TCP/IP networks, and provides troubleshooting information for both types of networks.

Intended Audience

The *NETWORK TROUBLESHOOTING GUIDE* is intended for users who are fairly new to network management and troubleshooting, but who have some VMS or ULTRIX system management experience.

Document Structure

The *NETWORK TROUBLESHOOTING GUIDE* consists of the following chapters and appendixes:

- The Preface provides introductory information, including the required skills, resources, and privileges needed to troubleshoot network problems.
- Chapter 1 reviews basic network concepts, including information related to DECnet and TCP/IP networks.
- Chapter 2 explains the structured approach to network troubleshooting.
- Chapter 3 gives an overview of network troubleshooting tools.
- Chapter 4 describes resources and routine procedures for solving network problems.
- Chapter 5 provides detailed information for solving specific network problems.
- Appendix A provides information on configuring Ethernet networks.
- Appendix B explains how to obtain Requests for Comments (RFC) documents.

Required Skills

Before troubleshooting network problems, you should be familiar with VMS or ULTRIX system management. You should understand the basic components of a data communications network, and have experience with network management tools appropriate to your network and operating system. For example, if your network runs DECnet networking software, you should be familiar with the Network Control Program (NCP), the software used to manage DECnet networks. If you have a TCP/IP internetwork, you should be familiar with ULTRIX commands such as arp, netstat, ping, syslog, and uerf.

Required Resources

This manual does not cover all the background information needed for effective troubleshooting. For more detailed information, see the system management and networking manuals of the VMS and ULTRIX documentation sets, the DECnet documentation appropriate to your operating system, as well as the documentation for the tools discussed in Chapter 3, Network Management and Troubleshooting Tools.

Some of the network management tools require specific system resources. See Chapter 3, Network Management and Troubleshooting Tools, for more information on these requirements.

Required Privileges

To troubleshoot network problems, you need the following privileges:

For VMS Systems:

- An account with system management level privileges on at least one, and preferably multiple systems in your network
- Privilege to read and write data files
- Privileges required to use various network management tools, as described in Chapter 3, Network Management and Troubleshooting Tools

For ULTRIX Systems:

- A user account
- Access to the superuser account

When you log in to the superuser account (also called the *root* account), the system processes your requests without performing the normal security checks. The superuser has complete access to any file and complete control of any process on the system. See the *ULTRIX-32 Introduction to System and Network Management* for more information on the superuser account.

Associated Documents

DECnet-VAX and VMS Troubleshooting

- *A Common Sense Guide to Network Management*
- *DECconnect System Planning and Configuration Guide*
- *DECelms Installation manual*
- *DECelms Reference manual*
- *DECelms Use manual*
- *DECmcc Management Station for ULTRIX Use manual*
- *Guide to DECnet-VAX Networking*
- *Guide to Terminal Server Manager*
- *LAN Traffic Monitor User's Guide*
- *NMCC/VAX ETHERnim User's Guide*
- *NMCC/DECnet Monitor User's Guide*

- *VMS Authorize Utility Manual*
- *VMS DCL Dictionary*
- *VMS LAT Control Program (LATCP) Manual*
- *VMS Network Control Program Manual*
- *VMS Networking Manual*

DIGITAL Network Architecture (DNA) Specifications

- *DECnet DIGITAL Network Architecture General Description*
- *DIGITAL Data Communications Message Protocol Functional Specification*
- *Network Services Protocol Functional Specification*
- *Maintenance Operation Protocol Functional Specification*
- *Data Access Protocol Functional Specification*
- *Routing Layer Functional Specification*
- *DNA Session Control Functional Specification*
- *DNA Phase IV Network Management Functional Specification*
- *Ethernet Node Product Architecture Specification*
- *Ethernet Data Link Functional Specification*

DECnet-ULTRIX and ULTRIX Troubleshooting

- *Introduction to System and Network Management*
- *Guide to System Configuration File Maintenance*
- *Guide to System Crash Recovery*
- *Guide to System Exercisers*
- *Guide to Networking*
- *ULTRIX Reference Pages*
- *DECnet-ULTRIX Release Notes*
- *DECnet-ULTRIX Installation and Checkout Procedures*

- *DECnet-ULTRIX User's and Programmer's Guide*
- *DECnet-ULTRIX Guide to the DECnet-Internet Gateway*
- *DECnet-ULTRIX Guide to Network Management*

TCP/IP Network Information

- *Requests for Comments (RFCs)*
Appendix B explains how to obtain RFCs.
- *Internetworking With TCP/IP: Principles, Protocols, and Architecture* by Douglas Comer (Prentice-Hall, 1988)
- *The Design and Implementation of the 4.3BSD UNIX Operating System* by Samuel J. Leffler, Marshall Kirk McKusick, Michael J. Karels, John S. Quarterman (Addison-Wesley, 1989)

Conventions Used in This Document

Convention	Meaning
Ctrl/x	Ctrl/x indicates that you hold down the Ctrl key while you press another key (indicated here by x).
User Input	In examples, user input is designated by red text in hardcopy documentation and bold text in online documentation.
<i>italic text</i>	Italic text emphasizes important information, indicates variables, references titles of other manuals, and indicates new terms that can be found in the glossary.
UPPERCASE in VMS and DECnet examples	In VMS and DECnet examples, words in uppercase indicate a command, the name of a file, the name of a file protection code, or an abbreviation for a system privilege.
UPPERCASE in ULTRIX examples	The ULTRIX operating system differentiates between lowercase and uppercase characters in commands. Enter uppercase characters only where specifically indicated by an example or a syntax line.
lowercase	In VMS and DECnet format descriptions, words in lowercase indicate parameters or arguments to be specified by the user.
%	The percent sign (%) is the default user prompt in multiuser mode for the ULTRIX operating system
#	The number sign (#) is the default superuser prompt for the ULTRIX operating system.

Overview of Data Communications Networks

This chapter provides a review of the basic concepts of data communications networks. It also provides an overview of two specific types of data communications networks: DECnet and TCP/IP (or Internet).

A firm understanding of these basic data communications concepts is critical to your understanding of networks in general, and to your work in network troubleshooting.

What Is a Data Communications Network?

In its simplest form, a data communications network (or computer network) is two or more computers set up to exchange information or data. In its most complex form, a computer network can consist of many different types of computers set up to exchange information and data throughout the world.

What Are the Components?

A computer network consists of four basic components:

- Computers

Computers consist of the CPU, memory, auxiliary storage, communications devices such as Ethernet controllers, and terminals. In a wide area network (WAN), the term **data terminal equipment** or **DTE** refers to the computer components. In a local area network (LAN), computers are known as nodes or hosts.

- Software

The software that runs on computers includes operating system software, such as VMS, ULTRIX, or MS-DOS; communications software such as DECnet or TCP/IP, or both; and application software, such as programming languages, text editors, and so forth.

- Data communications connections

Data communications connections provide the interfaces to the network for computers (or DTEs). In wide area networks (WANs), modems and multiplexors provide these interfaces, and this equipment is collectively known as **data communications equipment** or **DCE**. In local area networks (LANs), equipment such as transceivers provide these interfaces, and this equipment is known as **media access units** or **MAUs**.

Overview of Data Communications Networks

What Is a Data Communications Network?

- Data communications media
Data communications media connect the network components together, providing the **communications channel** (or path) through which signals can be sent to the computers or DTEs. Communications media come in several different forms including twisted-pair wire, coaxial cable, fiber-optic cable, microwave, and satellite.

Figure 1–1 illustrates one way that you can connect DTEs together using DCEs in a wide area network.

Figure 1–1 Point-to-Point Connections

Figure 1–2 is an example of how you can connect DTEs directly to an Ethernet cable in a local area network.

Figure 1–2 Ethernet Connections

What Does a Data Communications Network Look Like?

The way you connect the network components together can yield very different network configurations or topologies. **Topology** refers to the physical and logical location of components in a network. **Physical location** refers to the place a network device is stationed. For example, network devices can be located in a computer lab, a user's office, an office communications cabinet, or on a factory floor. **Logical location** refers to the functional interconnections between devices. For example, a node that is logically adjacent to another node can be physically located on another floor of a building. Functionally, the adjacent node is the closest node in terms of network **hops**, the distance between two directly connected nodes.

This section discusses the ways computer devices are physically connected to form a network, and gives brief explanations of star, ring, bus, tree, and mesh topologies. This section also addresses the ways networks are connected to form local area and wide area networks.

Overview of Data Communications Networks

What Is a Data Communications Network?

Star Topology

In the star topology show in Figure 1–3, each computer in the network is connected to a central computer through a data circuit. The central computer knows the paths to all the other computers, and processes all messages destined for the computers connected to it. Because all computers in a star topology connect to a central computer, the central computer presents the opportunity for a single point of failure on the network.

An example of a star topology is a wide area network (WAN) in which all computers connect to a computer at a central office site.

Figure 1–3 Star Topology

Ring Topology

A ring or loop topology, as shown in Figure 1–4, consists of computers arranged in a closed loop. Messages destined for another computer in the network pass from the sending computer to each computer along the loop until they reach their destinations. **Fiber Distributed Data Interface (FDDI)** and **token ring** networks are examples of ring topologies.

Overview of Data Communications Networks

What Is a Data Communications Network?

Figure 1–4 Ring Topology

Bus Topology

In a bus topology, all the computers are connected to a single cable that runs the length of the network, as shown in Figure 1–5. Every computer has access to every other computer on the network directly through the transmission media. An Ethernet network is an example of a bus topology.

Figure 1–5 Bus Topology

Mesh Topology

A mesh topology builds upon the topologies described thus far by ensuring that no single point of failure exists on the network. A mesh network accomplishes this through the use of multiple paths to each computer system, as shown in Figure 1–6. A DECnet wide area network configuration can be an example of a mesh topology.

Overview of Data Communications Networks

What Is a Data Communications Network?

Figure 1–6 Mesh Topology

Tree Topology

The tree topology provides a hierarchical approach to network configuration. A central, controlling computer manages the various groups of computers.

Figure 1–7 illustrates a simple tree-oriented network. A network that uses the IEEE 802.1 spanning tree algorithm for bridges interconnects the LANs in a tree topology.

Figure 1–7 Tree Topology

Local Area Network Topology

A **local area network (LAN)** is a data communications network that spans a physically limited area. A LAN is usually owned by the organization it services (that is, not provided by a common carrier) and provides high-bandwidth communication over inexpensive media. In a LAN, any or all of the topologies described thus far can exist.

Overview of Data Communications Networks

What Is a Data Communications Network?

LANs allow multiple protocols to share the same media. A **protocol** governs the operation of a communications link, and consists of a set of messages with specific formats and the rules for exchanging the messages.

For example, on an Ethernet, DECnet and TCP/IP can use the same Ethernet cable. The number of protocols running on a LAN depends on the needs of its users. For example, the LAN might be set up to allow some computer systems to run DECnet, and communicate only with DECnet nodes. Others might run TCP/IP and communicate only with TCP/IP hosts. If this type of LAN configuration used application gateway software, such as the VMS/ULTRIX Connection software, the DECnet nodes and TCP/IP hosts on the same LAN would have some ability to communicate.

In addition, computer systems on a LAN can run both DECnet and TCP/IP, enabling them to be both a DECnet node and a TCP/IP host. Such systems can communicate with both DECnet nodes and TCP/IP hosts.

Section 1.2.7.1 provides further LAN information and figures showing various DECnet LANs.

Wide Area Network Topology

A **wide area network (WAN)** covers a much larger geographic area than a LAN. Whereas a LAN usually spans an office building or campus, a WAN connects computer systems in distant locations, and around the world using common carriers and satellite links to transmit data. WAN networks are predominantly star, mesh, and tree topologies.

Section 1.2.7.2 provides further WAN information and figures showing various WAN configurations.

What Is a DECnet Network?

A DECnet network consists of two or more computing systems linked to exchange information and share resources. DECnet is the collective name for the family of communications products (software and hardware) that allow operating systems to participate in a network. DECnet networking software is available for a variety of operating systems including VMS (which uses DECnet-VAX), ULTRIX (which uses DECnet-ULTRIX), and MS-DOS (which uses DECnet-DOS). Many features of the DECnet networking software are common across these operating systems. For simplicity, examples in this manual focus primarily on DECnet-VAX and its operation with the VMS operating system.

An operating system uses its DECnet networking software to become part of a DECnet network. As a part of a DECnet network, a computing system can communicate with any other system on the network that uses DECnet software.

Each system on a DECnet network is called a **node**. DECnet networks are known as **peer-to-peer** networks because all systems participate in the network as peers or equals. Systems can communicate with each other without going through a central or master system. Any system in the network can communicate with any other system in the network, not merely with those systems to which it is directly attached. Depending on the network configuration and security, network users can gain access to software that does not exist on their local systems, and can communicate freely over the entire network.

Overview of Data Communications Networks

What Is a DECnet Network?

A DECnet network links computers into flexible configurations to exchange information, share resources, and perform distributed processing. DECnet distributed processing allows systems to be placed at locations where they are required while still having access to the facilities of other widely dispersed systems. Access to the network is available wherever it is needed, for example: executive offices, factory floors, laboratories, or field locations. Information can be exchanged efficiently between all parts of an organization or institution in a stable, integrated networking environment. An entire organization can be connected into a single unit by means of the network.

What Does a DECnet Network Look Like?

DECnet allows users to plan computer networks ranging from a few workstations linked together in one room to a very large network of computers distributed around the world. The DECnet network is designed to permit growth without disruption. The network can grow from a minimum of two nodes to a maximum of 64,449 nodes. DECnet configurations are flexible and can be expanded easily. Nodes can be located wherever required. Individual nodes can be added or relocated without impact on existing nodes or interruption of network operation.

An **area** is a logical grouping of nodes. A DECnet network can have a maximum of 63 areas, each containing a maximum of 1,024 nodes; an optimum number is approximately 300 to 500 nodes, depending on the network topology. In a multiple-area network, the network manager groups nodes into separate areas, with each area functioning as a subnetwork. Nodes in any area can communicate with nodes in other areas.

DECnet supports many different kinds of network topologies. Nodes located in a building or a complex of buildings can be connected in a LAN. The network can be expanded to include nodes at more geographically dispersed locations, connected in a WAN. In addition, systems on a DECnet network can use other Digital communications products to communicate with certain non-DECnet systems and networks in an integrated network environment. Section 1.2.7 provides more information and figures explaining the network topologies that DECnet supports.

What Communications Media Does DECnet Use?

Nodes in a DECnet network can be linked by various types of data transmission media. Local area network configurations are created using the various forms of Ethernet. Each system or device is connected to the Ethernet transmission media by a single line. Ethernet provides a single, shared network channel to which all nodes have equal access. Because Ethernet is a multi-access device, new nodes can be added without affecting existing nodes on the Ethernet.

Wide area networks use dedicated lines, telephone lines, microwave and satellite links, and fiber-optic links. Telephone lines may be leased to provide for permanent connections, or may be used as dialup lines for specific periods of time. Communication over analog telephone lines normally involves the use of **modems** at each end of the connection to perform conversion between the digital signals used by the computer and the analog signals used on the telephone line. Digital circuits require the use of a **data service unit (DSU)** or **channel service unit (CSU)** to perform signal conversion.

For a microwave link, a message is converted into microwave signals at the transmitting site and reconverted at the receiving location, which can be some distance away. Satellite links are generally used for long-distance communication, such as transoceanic communication.

Overview of Data Communications Networks

What Is a DECnet Network?

What is Ethernet/802.3?

Digital uses the IEEE Ethernet/802.3 standard for its Ethernet products. Understanding the Ethernet/802.3 architecture is critical for successful troubleshooting on Ethernet/802.3 LANs. This section gives some background on Ethernet/802.3.

An Ethernet local area network provides a communication facility for high-speed data exchange among computers and other digital devices located within a medium-sized geographic area. Ethernet is a network bus consisting of coaxial cable that uses **carrier-sense multiple access with collision detection (CSMA/CD)** for channel access and channel contention. Ethernet operates at the physical and data link layers.

The primary characteristics of the physical layer are as follows:

- Data rate of 10 megabits per second
- Maximum station separation of 2.8 kilometers
- Maximum number of stations equal to 1,024
- Use of coaxial cable
- Support of a branching nonrooted tree topology

The physical layer performs data encoding and decoding, and channel access including bit transmission, carrier sense, and collision detection.

The data link layer performs data encapsulation, including framing, addressing, and error detection; and link management, including carrier sense and collision detection. The data link layer is built on the physical layer and relies on the physical layer for various services, including transmitting and receiving data on communications links, and controlling operation of the link.

Ethernet supports a data transmission rate of 10 million bits per second in a limited area. The maximum limit on the distance between any two nodes on Ethernet is 2.8 kilometers (1.74 miles). A single Ethernet segment can support up to 1,024 nodes, using a maximum of 100 taps per Ethernet segment. Multiple Ethernet segments can be connected using bridges, forming an extended Ethernet. An extended Ethernet can support up to 8,000 nodes.

Ethernet is available in coaxial, shielded and unshielded twisted pair, fiber-optic cable, and microwave. For certain environments (for example, in an office or other area where personal computers and workstations are located), a thin, flexible cable called ThinWire Ethernet cable can be used. Fiber-optic cable can be used to connect Ethernet segments over short distances (less than 10 kilometers).

Ethernet allows multiple protocols to run on the LAN simultaneously. This guide covers network problems involving several protocols including LAT, DECnet, MOP, TCP/IP, ARP, and bridge protocols.

How Do Systems Communicate over DECnet?

Every node in a DECnet network has a unique name and address. Nodes in the network are connected by lines over which circuits operate (see Figure 1–8). A **line** is a physical path over which data passes from one node to another in the network. (The path may be over a cable or telephone line, or possibly a microwave or satellite link.) You can think of a **circuit** as a higher-level logical connection that operates over the physical connection. The circuit is the communications data path that carries information from one node to another. All

Overview of Data Communications Networks What Is a DECnet Network?

input and output (I/O) activity between nodes occurs over circuits. Multiple users can use each circuit. A node can be configured to have active circuits operating over a number of lines that connect the node to the other nodes in the network.

Figure 1–8 Network Nodes, Circuits, and Lines

A computing system can run many different processes and programs. For two processes to communicate with each other, they must have a way to establish contact and exchange data. DECnet permits computer processes running on the same or different nodes to communicate with each other over logical links. A **logical link** connects two processes and carries a stream of two-way communications traffic between the processes over one or more circuits.

The process or program to which a logical link is connected is called an **object**. On a VMS node, some objects are DECnet–VAX system programs (for example, the MAIL object); other objects can be user-written programs. For two programs to communicate over the network, the program on the local node must establish a logical link with an object on a remote node.

How Does DECnet Route Messages?

The process of directing a data message from a source node to a destination node is called **routing**. The route the data travels over the circuits in the network is called the **path**.

Messages can be exchanged between any two nodes in the DECnet network, even if the nodes are not directly connected to each other. For nodes that are not directly connected to be able to communicate, an intervening node along the data path must forward the data received from the source to the destination. Intervening nodes that receive data and forward it to another node are known as routing nodes (or **routers**). Nodes that do not forward data are called **end nodes**. Both routers and end nodes can send messages to and receive messages from other nodes, but only the router can forward messages on behalf of another node. A router can have more than one active circuit connecting it to the network. End nodes can have multiple circuits and receive data on all of those circuits, but can only transmit on one circuit.

DECnet supports routing within each DECnet area and a second, higher level of routing that links the areas, resulting in less routing traffic throughout the network. Nodes that perform routing within a single area are referred to as **level 1 routers**; nodes that perform routing between areas as well as within their own area are called **level 2 routers** (or **area routers**).

A router maintains **routing tables**, an information database about the availability of paths to the destination node and keeps it up-to-date by regularly exchanging routing information with other routers. The routing information includes the cost and the number of hops involved in sending data on a path to a destination node. The **circuit cost** is a number that the network manager assigns to a circuit between two nodes; the **path cost** is the sum of the circuit costs along the path to a given node. A **hop** on a DECnet network is the distance between two directly connected nodes; the **path length** is the number of hops along the path between two nodes.

Overview of Data Communications Networks

What Is a DECnet Network?

The router uses current information from its database to choose a data path through the network. The router determines the path to the destination based on the least cost. By changing the cost of a circuit, the manager of a network node can affect the flow of data through the network.

DECnet performs **adaptive routing** — that is, routing that adapts to changing conditions in the network. DECnet selects the best or **least cost path** currently available from the source to the destination. If network conditions change and the current least cost path becomes unavailable, DECnet redirects the data over the next best alternative path. DECnet automatically reroutes messages if a circuit becomes disabled or a lower-cost path becomes available.

Because adaptive routing in a DECnet network permits messages to be routed over the most cost-effective path currently available, a general user of the network need not be concerned with the path to the destination. Users need only specify the name or address of the remote node with which they wish to communicate.

How Is DECnet Software Structured?

DECnet software design is based on the DIGITAL Network Architecture (DNA), which follows industry standards. The structured design permits a DECnet network to be extended easily and to incorporate new developments in data communications. DECnet nodes can communicate with any system that supports the DECnet protocols.

DNA specifications govern the interrelationship of the components that make up the DECnet software. The specific functional boundaries between DECnet software components are structured as a hierarchical set of layers. Each DNA layer is a client of the next lower layer and does not function independently.

DNA specifies the functional layers in which DECnet software is arranged and the communications protocols through which the corresponding layers at different nodes communicate with each other.

Figure 1–9 illustrates the DNA layers and the related DNA protocols that provide DECnet network functions at each layer. For a complete description of DNA, see the DNA specifications. For a list of the DNA specifications, see the Associated Documents section in the Preface of this manual.

Overview of Data Communications Networks

What Is a DECnet Network?

Figure 1–9 DECnet Network Architecture (DNA) Layers and Protocols

Table 1–1 describes the functions of each of the DNA layers.

Table 1–1 DECnet Layers

DECnet Layer	Function
User	Contains most user-supplied functions and NCP.
Network application	Provides generic services to the user layer, including remote terminal access, remote file access and transfer, and resource-managing programs.
Session control	Defines the system-dependent aspects of logical link communication, including name-to-address translation, process addressing, and, in some systems, process activation and access control.
End communication (or Transport)	Provides for the system-independent aspects of creating and managing logical links for network users, including data flow control, end-to-end error control, and segmentation and assembly of user messages. In DECnet Phase V, this layer is known as the transport layer.
Routing (or Network)	Routes user data from its source to its destination, and provides congestion control and packet lifetime control. In DECnet Phase V, this layer is known as the network layer.

(continued on next page)

Overview of Data Communications Networks

What Is a DECnet Network?

Table 1–1 (Cont.) DECnet Layers

DECnet Layer	Function
Data link	Creates a communications path between adjacent nodes, and ensures integrity of the data transferred across this path. The data link layer includes multiple protocols including DDCMP, Ethernet, and X.25.
Physical link	Manages the physical transmission of data over a channel, including monitoring channel signals, clocking the channel, handling hardware interrupts, and informing the data link layer when a transmission is complete. The physical layer includes parts of device drivers, and hardware, such as any DCE equipment.

What Systems Can Communicate over the Network?

Digital communications software allows Digital systems to communicate with each other and with some non-Digital systems.

A VMS system, installed as a DECnet–VAX node on the network, can communicate directly with any other VMS systems and with any other Digital system connected to the same network. All members of the VAX family of processors on which VMS systems are running can be linked.

A VMS system can also communicate with any other Digital operating system on the network. For example, a VMS system running DECnet–VAX software can communicate with an ULTRIX system running DECnet–ULTRIX software, an RSX system running DECnet–RSX software, and a Professional 300–series system that uses PRO/DECnet software.

Several types of personal computers can join the DECnet network as end nodes. For example, certain IBM personal computers and IBM-compatible personal computers can participate as DECnet end nodes by means of DECnet–DOS or DECnet for OS/2 software.

DECnet–VAX nodes can use VAX PSI software to directly access a packet-switching network (such as Telenet or TYMNET) or can access the packet-switching network by means of an X.25 communications server. Nodes and terminals connected to a packet-switching network can use that network to communicate with each other. Packet-switching networks are often used for communication over very long distances involving common carriers and satellite links.

Through special interconnect products, such as gateways and emulators, DECnet nodes can communicate with non-Digital systems and networks. The DECnet /SNA gateway permits a DECnet network to connect to an IBM System Network Architecture (SNA) network.

What Network Environments Does DECnet Support?

DECnet networks support a variety of network connections, permitting computers to be linked in flexible configurations. LANs and WANs run a variety of Digital and non-Digital products, and can be combined to form integrated networks that allow different types of systems to communicate and share information and resources.

The LANs and WANs that make up a Digital network can be configured in most of the common topologies discussed in Section 1.1.2, including star, bus, ring, tree, and mesh topologies.

Overview of Data Communications Networks

What Is a DECnet Network?

Local Area Networks

The Digital local area network uses Ethernet to create local area networks. Local area networks can be configured in a variety of arrangements. Two Ethernet segments can be connected by means of a **bridge**, a relay that controls network traffic between the Ethernets it connects. Use of a bridge can extend a local area network beyond the distance limitation imposed on a single Ethernet. The METROWAVE bridge allows you to connect geographically separated 802.3 Ethernet LANs within a metropolitan or campus environment when cable is neither feasible nor economical. Routers can also be used to connect two Ethernets. In addition, routing nodes on an Ethernet can be connected to wide area network nodes to form a large, integrated network.

You can connect individual systems directly to an Ethernet, or you can connect several systems to a local network interconnect device, or *DELNI*, which can then connect to the Ethernet. The DELNI serves as a concentrator, grouping systems together.

Individual users can also gain access to the nodes in a local area network through a **terminal server**. A user at a terminal connected to the terminal server can access any service node that implements the local area transport (LAT) protocol and is known to the server. A user logged in to a node by means of a terminal server can perform the same functions as a user logged in to a terminal directly connected by an asynchronous line to the node.

Figure 1–10 illustrates a small Ethernet configuration of four nodes. Three MicroVAX-based end nodes and one VMS router are connected to the Ethernet.

Figure 1–10 Small Local Area Network Configuration

Figure 1–11 shows a larger local area network configuration in which two Ethernets are connected by a LAN bridge. Various kinds of operating systems, including the nodes in a Local Area VAXcluster, are connected directly to the Ethernet. In the figure, a group of small systems is connected to the Ethernet by means of a DELNI device. Individual terminal users can gain access to Ethernet nodes through a terminal server.

Wide Area Networks

A wide area network provides for communication over broader geographic areas. DECnet supports long-distance communication with systems located anywhere in the world. A wide variety of communications media can be used: examples include dedicated, leased and dialup lines, and microwave and satellite links. Nodes in a wide area network can be connected by point-to-point connections or through packet-switching networks.

Overview of Data Communications Networks

What Is a DECnet Network?

DECnet-VAX offers comprehensive wide area network support and long-haul connectivity over point-to-point connections. **Point-to-point connections**, which use the Digital Data Communications Message Protocol (DDCMP), are synchronous or asynchronous. **Synchronous** devices provide high-speed connections over dedicated lines or telephone lines (using modems). **Asynchronous** devices provide low-speed, low-cost connections over terminal lines that are switched on for network use either permanently (a static connection) or temporarily (a dynamic connection). For example, a user on a MicroVAX can configure a dialup line (a telephone line) to another computer as a dynamic asynchronous DECnet line for the duration of a telephone call.

Figure 1-11 Large Local Area Network Configuration

Overview of Data Communications Networks What Is a DECnet Network?

Figure 1-12 illustrates the different kinds of DDCMP connections.

Figure 1-12 DDCMP Connections

DECnet-VAX supports worldwide communications through packet-switching networks and gateways. A DECnet-VAX node can be connected to a packet-switching network (either directly using VAX PSI software or through an X.25 communications server) to establish communication with a remote DECnet node. Packet-switching networks, such as TYMNET and Telenet, provide communication services between nodes on the same or different networks, often in widely dispersed geographic areas connected by satellite links. A DECnet-VAX node can use a DECnet/SNA gateway to communicate with IBM systems in an SNA network.

Figure 1-13 shows various wide area network connections. The figure illustrates the use of a DDCMP synchronous line to connect two VMS nodes at different locations (Boston and New York). The nodes in Boston and London are both configured to use VAX PSI software, permitting the nodes to be connected directly to an X.25 packet-switching network. Through the packet-switching network, Boston can communicate with London. The figure also shows how the VMS nodes located in Boston, London, and New York can communicate with IBM systems on an SNA network by means of a DECnet/SNA gateway.

Overview of Data Communications Networks

What Is a DECnet Network?

Figure 1–13 Wide Area Network Connections

DECnet local area networks and wide area networks can be combined to provide comprehensive network support. Wide area network connections connect individual local area networks, and can provide access to non-Digital systems.

Figure 1–14 is an example of a large DECnet configuration that illustrates various ways VMS operating systems on VAX processors can be connected to the network. The figure indicates whether a particular VMS node is a router or an end node. The figure also shows a terminal server connected to the Ethernet. Individual terminal users can log in to any node on the extended Ethernet by means of the terminal server, provided the node has been defined on the server.

The DECnet configuration shown includes two Ethernet cables linked by a bridge. Two clusters of VMS nodes are attached to the Ethernet:

- A Local Area VAXcluster that includes three large VMS systems
- A Local Area VAXcluster that includes four other VMS systems (linked by their own ThinWire Ethernet, and isolated by a bridge)

In Figure 1–14, VMS routers and end nodes that are not members of a cluster are also connected to the Ethernet. The network diagram shows how DECnet–VAX point-to-point connections are integrated with the Ethernet multiaccess configuration:

- A synchronous DDCMP point-to-point line provides a connection to additional DECnet nodes at a remote location.
- Asynchronous DDCMP lines provide permanent (static) and temporary (dynamic) connections between VMS nodes in the network.

Figure 1–14 Large Integrated DECnet Configuration

In the figure, VMS nodes on either Ethernet can communicate with nodes at distant locations (through a packet-switching network) and can access an IBM SNA network by means of a gateway node.

What Is an Internet Network?

An **internetwork** or **internet** is a communications technology that enables you to interconnect physical computer networks so that you can exchange data among them. The interconnected networks often consist of widely varying types of hardware. Yet, once part of an internet, the networks function as a single, coordinated whole, regardless of the underlying types of hardware each network uses.

Overview of Data Communications Networks

What Is an Internet Network?

One of the largest internet implementations consists of participants from major research institutions, universities, and government labs, including the National Science Foundation (NSF) and the NSFnet regional organizations. This collection of computing networks is known as the TCP/IP Internet, or simply, the Internet.

The Internet is not a commercial product but rather, a large project in support of research. All work, proposals, and protocol standards appear in a series of technical papers called **Requests for Comments (RFCs)**. The **Network Information Center (NIC)**, the central authority for the Internet, distributes the RFCs to the Internet community using electronic or postal service mail.

You can obtain RFCs electronically if your host is part of the Internet. Appendix B gives instructions for requesting RFCs.

Note

This manual discusses internets in general, and the TCP/IP Internet in particular. To distinguish between the general and the specific, references to internets in general are in lowercase characters. References to the TCP/IP Internet are capitalized.

What Does an Internet Look Like?

Internet protocols allow users to plan computer networks of any size and arrangement, from a few workstations linked together in one room to a very large network of computers distributed around the world. Internet technology permits growth without disruption, from a minimum of two computer systems (or **hosts**) on a network to several million hosts.

Internet configurations are flexible and can be expanded easily. Hosts can be located wherever required. Individual hosts can be added or relocated without impact on existing hosts or interruption of network operation.

Internet protocols do not require any particular physical network media; they can be implemented on any network. Hosts located in a building or a complex of buildings can be connected in an Ethernet LAN or a token ring LAN. The network can be expanded to include hosts at more geographically dispersed locations, connected in a WAN. Internet networks can be configured using many types of communications media, and in any type of topology, including star, bus, ring, tree, and mesh.

Because of its inherent flexibility, an internet can consist of nearly any type of computer system running nearly any type of operating system, provided that the system has a communication device and can implement the internet protocols. For example, an internet can consist of workstations running ULTRIX, personal computers running MS-DOS, and minicomputers running VMS.

Figure 1-15 shows a simple internet configuration consisting of two physical networks connected by an internet protocol (IP) router. An **IP router** (also commonly called an **IP gateway**) is a host that connects to two or more networks. Each network consists of many hosts. The IP router knows how to reach all the hosts on both Network 1 and on Network 2, and controls all communication between the two networks.

Overview of Data Communications Networks

What Is an Internet Network?

Figure 1–15 Two Networks Interconnected to Form an Internet

Figure 1–16 shows a more complex internet configuration. In this configuration, two IP routers link three separate networks to create an internet. As with the Figure 1–15, each IP router in this configuration knows only about the hosts on the networks to which the router is directly attached. For information about how communication occurs on a network such as this, see Section 1.3.6.3, which discusses routers and communication between networks in more detail.

Figure 1–16 Multiple Networks Interconnected to Form an Internet

What Is TCP/IP and How Is It Structured?

This manual covers the leading internet technology known as **TCP/IP**, which is used as the basis for the Internet. TCP/IP is named after the two main protocols used on internets: **Transmission Control Protocol** and **Internet Protocol**.

TCP/IP network software is loosely based on the four-layer Internet Protocol model shown in Figure 1–17. This structured design allows networks to be easily extended, and to incorporate new developments in data communications.

Figure 1–17 Internet Protocol Model

The Internet Protocol model specifies the functional layers of TCP/IP software on each host, and the communications protocols through which the corresponding layers at different hosts communicate with each other. Each protocol is a set of messages with specific formats and rules for exchanging the messages. These protocols govern the operation of a communications link without regard to hardware.

Overview of Data Communications Networks

What Is an Internet Network?

Table 1–2 describes the functions of each of the TCP/IP layers.

Table 1–2 TCP/IP Layers

TCP/IP Layer	Function
Application	Sends data to, or receives data from the transport layer. The type of data the application layer sends to the transport layer is messages or streams.
Transport	Provides communication between application programs. This communication is also known as end-to-end communication. This layer, implementing the transmission control protocol (TCP), also provides reliable transport, ensuring that data arrives without error and in the correct sequence. The transport layer divides the streams of data received from the application layer into transport protocol packets and sends the packets to the internet layer.
Internet	Provides machine-to-machine communications, specifying the format of the internet packets known as IP datagrams , and forwarding the IP datagrams to the network access layer. The internet layer provides unreliable, connectionless, best-effort data transfer.
Network access	Accepts IP datagrams, and transmits them over a network. Operations at this layer use physical and data link functions, and involve the transfer of network-specific frames of data.

For a complete description of the Internet Protocol model, see RFC 791 *Internet Protocol*. See Section 1.3.7.1 for a brief overview of IP and Section 1.3.8.1 for an overview of TCP.

What Are Addressing, Naming, and Routing?

In understanding how an internet operates, you must understand how the network locates and identifies hosts on the network, and how the network sends internet packets known as **IP datagrams** from one host to another. These functions are known as addressing, naming, and routing, respectively.

- **Addressing** is how the network defines the locations of hosts
- **Naming** is how hosts are identified
- **Routing** is how IP datagrams are sent from one host to another

The following sections focus on conventions for identifying and defining the location of hosts on an internet, and on the ways hosts use this information to communicate with each other.

Addressing

Each computer connected to an internet is identified by at least one **internet address**. Strictly speaking, an internet address does not identify the host itself; it identifies the host's connection to the network. Because a host can have multiple connections to the network, it can have multiple internet addresses. In addition to the internet address, each host can have a unique physical (or hardware) address. This section discusses only the internet address. Section 1.3.9.1 discusses the relationship between the internet and physical addresses.

Overview of Data Communications Networks

What Is an Internet Network?

A host on an internet network can have more than one connection to the network; thus, more than one internet address. For example, because IP routers interconnect multiple networks, each IP router has a separate internet address for each of its network connections.

An internet address consists of two parts, as shown in Figure 1–18: the **network number** of the host, and the **host number**.

Figure 1–18 Parts of an Internet Address

The network part of the internet address encodes the network class, as well as the network address. (Section 1.3.4.2 explains network classes.) The host part of the internet address identifies a unique host on the network.

The network number must be the same for all hosts connected to the same network. The host number for each host must be unique to the network to which it is connected. No two hosts on the same network can have the same host number. On the Internet, the Network Information Center (NIC) assigns network numbers and ensures that no two networks connected to the Internet have the same network number. The local network administrator ensures that host numbers on the local network are unique.

An internet address uses 32 bits organized into four fields (called **octets**) to specify the network number and host number for a computer on an internet. Figure 1–19 shows the 32 bits and their arrangement into octets.

Figure 1–19 Internet Address Fields

Common Internet Address Notation

At the machine level, the internet address is in bits, however these addresses are commonly translated into a more readable form known as **common internet address notation** or **dotted decimal notation**. Figure 1–20 shows the correspondence between a 32-bit address and the common internet address notation for the same address.

Figure 1–20 Common Internet Address Notation

Overview of Data Communications Networks

What Is an Internet Network?

Network Classes

On the Internet, the NIC (not the local network administrator) categorizes networks into classes, based on the number of hosts participating in each network. Network class information is included in the network number part of the internet address. The class of a network determines how the four fields of the internet address are used. The three main internet network classes are A, B, and C. Class A networks (such as the Internet) are the largest, and Class C are the smallest, as shown in Table 1-3.

Overview of Data Communications Networks

What Is an Internet Network?

Table 1–3 Sizes of Internet Network Classes

Network Class	Size
Class A	More than 65,536 hosts
Class B	From 256 to 65,536 hosts
Class C	Less than 256 hosts

You can determine the network class of an internet address from the most significant bits.

If the most significant bit is 0, the network class is A

If the most significant bits are 1 and 0, the network class is B.

If the most significant bits are 1, 1, and 0, the network class is C.

Figure 1–21 shows how the internet address bits are assigned to the network and host numbers for the three classes of networks.

Figure 1–21 Internet Address Format for Network Classes

Only a few organizations meet the criteria for Class A networks. Thus, the network part of the internet address required to identify Class A networks can remain relatively small. Only one octet of the 32-bit address identifies the network part of a Class A network address, and this octet specifies both the network class and network number. The first bit of a Class A network is 0. Because the number of hosts in a Class A network can be large, Class A networks use three octets of the 32-bits to identify hosts.

Mid-size, Class B networks (those ranging from 256 to 65,536 hosts) are more numerous than Class A, and require a larger portion of the 32-bit address to identify the network address of hosts. In Class B networks, the first two octets identify the network, and the last two octets identify the host.

Class C networks are small (consisting of fewer than 256 hosts), however, they are more numerous than Class A or Class B networks. Thus, they require three octets to identify the network part of the internet address, but only one octet to identify the host part of the address.

Overview of Data Communications Networks

What Is an Internet Network?

In specifying the network part of an internet address, each class is allowed a range of numbers. Table 1–4 shows the ranges for network classes A, B, C, and D, and gives an example of an internet address for each class.

Table 1–4 Address Ranges for Network Classes

Network Class	Address Range ¹	Internet Address Example
A	2.0.0.0–126.0.0.0	98.12.5.32
B	128.1.0.0–191.254.0.0	130.9.10.15
C	192.0.1–223.254.254	201.43.89.3
D ²	224.0.0.0–239.255.255.255	230.25.25.25

¹The numbers 1 and 127 are reserved for other purposes and are not used in identifying network classes. For example, on ULTRIX systems, the address 127.0.0.1 normally designates localhost, the software loopback interface.

²Used for internet multicasting.

Subnet Addressing

The development of subnet addressing helped solve problems that occurred due to the proliferation of workstations, personal computers, and small networks on the Internet. Before subnet addressing, each physical network (or LAN) required a unique network address. As the number of networks participating in the Internet increased, Internet administration became increasingly complex. Subnet addressing helped to simplify and distribute administration and communications on the Internet by allowing multiple physical networks to share a single network address.

Each physical network that shares a network address with other physical networks is known as a **subnet**. Because a single network address identifies all of the physical networks in a subnet collection, networks outside of the collection see only one network, where in fact there are multiple physical networks.

An internet network without subnets consists of two levels: the overall internet network, and the individual networks that make up the internet. When individual networks implement subnets, the internet network hierarchy expands to contain a third level, resulting in the network structure shown in Figure 1–22.

Overview of Data Communications Networks What Is an Internet Network?

Figure 1–22 Internet Network Hierarchies

Overview of Data Communications Networks

What Is an Internet Network?

Figure 1–23 shows an internet that does not use subnets. Each network (128.10, 130.14, and 130.12) is a separate internet network.

Figure 1–23 Internet Without Subnets

Figure 1–24 shows the same network with subnets. Hosts on other internet networks see only network address (140.5.n.n) for IP Router D, not the subnets. Remote hosts communicate with all of the hosts associated with subnets 130.12.n through IP Router D. The router sends data to the appropriate subnet, where the destination host recognizes its own address and picks up its data.

Figure 1–24 Internet with Subnets

Advantages of Subnet Addressing

Subnets improve internet operations and communications by clustering networks under one network number and reducing the total number of network addresses on the internet. Because fewer network addresses exist, the size of the routing tables that the IP routers maintain can remain smaller.

The reduction in routing table information is important for two reasons: First, because some routers would not have room to accommodate all the networks that would exist if subnets were not used, and second, because smaller routing tables can help improve IP router performance.

Subnets also have advantages for the local network manager. For example, the local network manager assigns subnet numbers, not the NIC (which assigns Internet network numbers). As a result, the local network manager has more flexibility when organizing network hosts and creating LANs to meet the needs of users. Without subnet addressing, the local network manager must obtain a new network number from the NIC each time the site requires a new LAN, and notify the NIC when removing a LAN.

Subnets and Internet Addresses

Use of subnet addresses requires a slight modification of the internet addressing scheme. Instead of the network number and host number associated with a standard internet address, a subnet address divides the 32-bit address into a network part and a local part. The network part identifies the internet network address. The local part identifies the physical network (subnet) and host number. Figure 1–25 shows the parts of an internet subnet address for a host on a Class B network.

Overview of Data Communications Networks

What Is an Internet Network?

Figure 1–25 Internet Subnet Address Parts

The subnet addressing structure in Figure 1–25 is the recommended way of setting up an internet subaddress, however, other methods exist. For more information on the setup variations for subnet addresses, see RFC 950, *Internet Standard Subnetting Procedure*.

Address Masks

Section 1.3.4.3.2 explained how subnet addresses modify the 32-bit internet address to specify a subnet number as well as a host number and primary network number. Because a subnet address looks the same as a standard internet address, hosts need to have a way to decode or interpret internet addresses to determine if a subnet address is in use.

The **address mask** (also called a **network mask**, **netmask**, or **subnet mask**) allows hosts to evaluate the destination address on incoming and outgoing datagrams to determine if the destination address contains a subnet number as well as a network number and host number.

Each host defines its address mask in its `/etc/rc.local` setup file. The address mask is a 32-bit number (four octet fields). Each of the 32 bits in the address mask correspond, one-to-one, with each of the 32 bits in the internet address.

When a bit in the address mask is on (binary 1), the corresponding bit position in the internet address is interpreted as part of the network or subnet address. When a bit in the address mask is off (binary 0), the corresponding bit position in the internet address is interpreted as part of the host address. For example, the decimal number 255 is 11111111 in binary notation, and means that the field should be interpreted as part of the network or subnet address.

Field 1 of the address mask is always 255 because it is interpreted as the network address, regardless of whether there are subnets or not. Field 4 is usually 0, so the system can interpret the host address. Fields 2 and 3 are usually 255 or 0, depending upon how you define the address mask. The way fields 2 and 3 are interpreted depends on the class of network.

For example, an address mask of 255.255.0.0 could be one of the following:

- The default address mask for a Class B network, with field 1 and field 2 defining the network number, and field 3 and field 4 defining the host number, as follows:

(1) network field (2) network field (3) host field (4) host field

- A Class A network with field 1 defining the network number, field 2 defining the subnet number, and field 3 and field 4 defining the host number, as follows:

(1) network field (2) subnet field (3) host field (4) host field

Overview of Data Communications Networks What Is an Internet Network?

Remember, only bits in the host portion of the internet address designate subnets. The way to determine the type of network an address mask designates is to look at the internet address, and check its address range. See Table 1–4 for the network classes for various network ranges. Table 1–5 shows the internet address fields available to designate subnets for each network class. An asterisk (*) indicates that the field is available for subnet designation.

Table 1–5 Internet Address Fields Available For Subnets

Network Class	Field 1	Field 2	Field 3	Field 4
A		*	*	*
B			*	*
C				*

In general, an entire 8-bit field is either on (255) or off (0). This makes it easier for users to distinguish among the network, subnet, and host portions of the address. However, values other than 255 and 0 can be used also. For example, if a Class C network implements subnets, it must use part of field 4 to designate both the subnet and the host.

Naming

Section 1.3.4 describes one way to identify network resources—the unique 32-bit address. However, the 32-bit address is difficult to use and remember, so users prefer to assign meaningful names for network resources.

In small networks it is fairly simple to assign names to each resource, and to keep track of the relationships between the names and the 32-bit addresses. However, the larger a network grows, the more difficult it is to assign unique names to resources, and to maintain the correct mapping information from the name to the 32-bit address.

In maintaining internet networks and troubleshooting problems, it is important to understand how naming works, and the specific software your network uses to keep track of name and address translations.

This section describes the internet naming convention known as the domain name system, and the services available for keeping track of the relationships between names and addresses.

Domain Name System

The **domain name system** is a tree-structured system for organizing host names for an entire internet. Naming information is organized into the following categories, each describing a more specific set of network hosts:

- Root
- Domain
- Subdomain (also known as a zone or branch)
- Leaf (the host name itself)

The **root** of the domain name system encompasses seven primary categories of internet hosts known as **domains**. Each domain serves a specific segment of the internet community. This section uses the Internet domain names, shown in Table 1–6, as examples.

Overview of Data Communications Networks

What Is an Internet Network?

Table 1–6 Internet Domains

Domain Name	Internet Group
COM	Commercial organizations
EDU	Educational institutions
GOV	Governmental institutions
MIL	Military groups
NET	Major network support centers
ORG	Organizations other than those above
ARPA	Temporary ARPANET domain

The Internet domain name system also allows countries other than the United States to participate in the Internet. Special domains, identified by international standard 2-letter country codes enable other countries to participate in the Internet.

The domain name system classifies each domain according to the various groups that participate in the domain. These groups are known as **subdomains**, **zones**, or **branches**. For example, the COM domain consists of commercial organizations such as Digital Equipment Corporation, Hewlett Packard Corporation, and other corporations. The EDU domain consists of educational institutions such as Purdue and Stanford Universities.

In the EDU domain, each subdomain is further classified according to the groups it comprises. The XYZ University (XYZ) subdomain consists of groups organized according to group functions. These subdomains within subdomains finally break down into individual hosts. Each host is a **leaf** in the domain name system.

Figure 1–26 shows the structure of the Internet domain name system, detailing some potential components of the EDU domain. Note that the EDU domain consists of many more subdomains than those listed here.

Figure 1–26 Internet Name Domains

In the domain name system, each domain maintains naming information for the levels directly below it. For example, the EDU domain is responsible (or authoritative) for all the subdomains and leafs below it.

However, if all naming information for the domain remained at the EDU domain level, naming service would become cumbersome. To provide more efficient naming services, the domain naming system allows domains to distribute authority for naming functions to lower levels in the domain. If the subdomains are large, the subdomains may also distribute some naming authority to lower levels.

For example, in the EDU domain, ABC and XYZ each have authority for all the naming for their own subdomains. In the XYZ subdomain, some authority is distributed to the next lower level, which includes GROUP1 and GROUP2. GROUP1 has naming authority for all the hosts in its group. Likewise, GROUP2 has naming authority for all of its hosts (or leafs).

Name Servers

A **name server** is software that maintains naming information and performs name-to-address and address-to-name translations. Sometimes the name server software runs on a dedicated processor, and the processor itself is also known as the name server.

To find out how to reach another host, the sending host obtains naming information about the receiving host. To do this, the sending host (known as a **client** in this case) sends a request through software known as a resolver. The **resolver** is client software that queries name servers for the naming information requested. Figure 1–27 illustrates the name server process.

Overview of Data Communications Networks

What Is an Internet Network?

Figure 1–27 Name Server Process

Name servers can be distributed throughout the network, providing efficient and reliable naming information. The top-level name servers for the domain name system shown in Figure 1–26 are **root servers**. Additional servers exist throughout the network, often (although not necessarily) corresponding to domain and subdomain levels.

Two of the most common name servers used in TCP/IP networks are the **Berkeley Internet Name Domain (BIND)** service, and **Yellow Pages (YP)**.) In addition, the `/etc/hosts` file on each host contains a list of known hosts on the internet, and information such as the hosts' names and internet addresses. The `/etc/hosts` file can provide some degree of name and address information, although the `/etc/hosts` file is not a name server.

Routing

Routing is the process computers use to direct an IP datagram from one host to another. Two kinds of routing can occur on internet networks: **direct** and **indirect**. All routing is based on the network and subnet number parts of the internet address, and use of IP routing tables. This section describes routing tables, and direct and indirect routing. It also provides examples showing how routing occurs on networks that use subnets.

IP Routing Tables

All hosts, including routers, use a table-driven method to route and deliver IP datagrams to their final destinations. Hosts maintain **routing tables** to keep track of the best routes to other networks in an internet. For each destination network, routing tables list the network number, and the internet address to use when sending datagrams to that network. If the destination network is the same as the one to which the host is attached, the host sends the datagram directly. If the destination network is different from the one to which the host is attached, the host consults its routing table to find the address to use to forward the datagram to its destination.

IP routers maintain primarily network and subnet information in their routing tables. Nonrouter hosts maintain primarily host-specific information in their routing tables. As a result, routing tables on each router can remain relatively small, and these relatively small routing tables can facilitate communication with hosts throughout an internet.

Overview of Data Communications Networks

What Is an Internet Network?

Direct Routing

Direct routing occurs when a computer sends a datagram to another computer on the same physical network. The sending host determines whether the destination host is on the same network by comparing the network part of the destination address to the network part of each address for which the host has an interface enabled.

The host's comparison includes applying an address mask to each address to obtain the network portion of the address. The address mask always includes the subnet bits, if the network uses subnets.

If the network address matches any of the networks to which the host is directly connected, the host sends the datagram over the appropriate interface directly to the designated host.

If the network part of the address does not match any of the networks to which the host is directly connected, the host continues with indirect routing, described in Section 1.3.6.3.

Indirect Routing

In cases where a datagram is destined for a host on another physical network, the sending host sends the datagram to an IP router on the same physical network. The IP router connects multiple physical networks, and can forward datagrams between any of the physical networks to which it is directly connected. The datagram continues to be sent from router to router until the datagram reaches a router that has direct access to the destination host.

Sometimes during this process, a router may receive a datagram from a host that suggests using a less than optimal route. In this case, the router sends an ICMP redirect message to the source host. The ICMP redirect message tells the source host of the more direct route to the destination host. The source host then updates its routing tables with this information.

The following steps describe indirect routing in more detail.

1. The host checks to see if its routing tables contain an explicit host-specific route to use to reach the destination host. If a host-specific route exists, the host uses that route, sending the datagram to the IP router defined in the host-specific route.

If the routing tables do not contain a host-specific route for the destination host, the host continues with step 2.

2. The host checks to see if its routing tables contain an entry for the destination network. If an entry exists, the host uses the route specified in the routing tables. The host sends the datagram to the IP router specified for the destination network.

If an entry does not exist, the host continues with step 3.

3. The host checks to see if the routing tables contain an entry for a default route. A default route is one that the host uses for any destination networks not listed in the host's routing tables.

If the routing tables specify a default route, the host uses that route, sending the datagram to an IP router that is responsible for routing to the default route.

Overview of Data Communications Networks

What Is an Internet Network?

If the routing tables do not specify a default route, the host returns an the error message, "network unreachable," to the user who sent this message.

Note

When a datagram reaches an IP router , the process begins again at step 1. The message continues to be sent from router to router until it reaches a router that is directly connected to the destination host's network.

Figure 1–28 illustrates the internet routing process.

Figure 1–28 Internet Routing Process

Routing In the Presence of Subnets

This section provides several examples to show how subnets affect internet routing. All examples are based on the configuration shown in Figure 1–29.

Incoming Datagrams to Subnets

The following example describes the routing process for an incoming datagram, originating from an internet network outside of 130.12, and bound for Host X (130.12.2.1), as shown in Figure 1–29.

1. The datagram arrives at IP Router D, the only known connection into network 130.12.

To the sending host, it is irrelevant whether the destination host is in a subnet or not. Information about the subnet concerns only IP Router D.

Overview of Data Communications Networks

What Is an Internet Network?

2. IP Router D checks the incoming datagram's destination address, and the router's subnet mask.
From this information, Router D determines that the datagram is for a subnet of 130.12.2.
3. IP Router D checks its routing tables for the correct route to network 130.12.2 and finds that interface 130.12.5.2 is the correct path to use.
4. IP Router D forwards the datagram through interface 130.12.5.2 to IP Router B.
5. IP Router B receives the datagram through interface 130.12.5.1 .
6. IP Router B checks its routing tables and finds that network 130.12.2 is a directly connected network.
7. Router B sends the the datagram through network interface address 130.12.2.3. The datagram travels directly to Host X (130.12.2.1).

Outgoing Datagrams from Subnets

The following example describes the routing process for an outgoing datagram, originating on Host X (130.12.2.1), and bound for an internet network outside of network 130.12, as shown in Figure 1–29.

1. Host X sends the datagram to IP Router B.
Host X knows that the destination host is not directly connected to subnet 130.12.2, because the network portion of the destination address does not match Host X's network.
2. IP Router B checks its routing tables and finds that the outgoing datagram is destined for a network outside of 130.12.
3. IP Router B forwards the datagram (through interface 130.12.5.1) to IP Router D (interface 130.12.4.2).
4. IP Router D checks its routing tables and forwards the datagram (through interface 140.5.n.n) toward its destination.

Datagrams from One Subnet to Another Related Subnet

The following describes the routing process within a set of subnets. It assumes a datagram originating from Host V (130.12.3.1) destined for Host Y (130.12.1.2), as shown in Figure 1–29.

1. Host V checks the destination address against Host V's subnet address mask and routing table
Host V finds that Host Y is not on the same physical network, but on a related subnet of network 130.12.
2. Host V sends the datagram to IP Router C (130.12.3.3).
3. IP Router C forwards the datagram (through interface 130.12.6.1) to IP Router D (interface 130.12.6.2).
4. IP Router D checks its routing tables to determine the path to get to network 130.12.1 and finds that the correct path is through network 130.12.4.
5. IP Router D sends the datagram (through interface 130.12.4.2) to IP Router A (interface 130.12.4.1).
6. IP Router A receives the datagram and checks the destination address, finding that the destination address is on subnet 130.12.1.

Overview of Data Communications Networks

What Is an Internet Network?

7. IP Router A forwards the datagram (through interface 130.12.1.3) using direct routing to Host Y (130.12.1.2).

Figure 1–29 Subnet Routing Examples

Network Protocols

Network protocols specify how hosts send data to other hosts on an internet. This section discusses the Internet Protocol (IP), routing protocols including the Routing Information Protocol (RIP), and the Internet Control Message Protocol (ICMP), an error reporting protocol.

Internet Protocol

Internet Protocol (IP), is the layer upon which TCP and application services build. IP determines how data is sent from one host to another on an internet, and specifies the format of internet packets known as IP datagrams.

The functions of IP are commonly described as follows:

- Unreliable, because IP may lose data, or may deliver it out of sequence, or may duplicate it
- Connectionless, because related packets are delivered independently of each other
- Best-effort, because IP does not fail unless resources are exhausted, or the underlying physical networks fail

Layers above IP provide reliable, connection-based service.

Routing Protocols

Routers use routing protocols to keep other routers informed of the best paths to any given network, and the availability of those paths. Included in this information is the address of the source router, and information on how far away the destination network is from that source.

The following types of routing protocols exist to address internet routing needs:

- **Interior Gateway Protocol (IGP)**

An IGP handles routing information for networks within an autonomous system. An **autonomous system** is a collection of routers and networks that fall under one administrative entity and cooperate closely to propagate network reachability and routing information among themselves.

Many IGPs exist, but some of the more common IGPs are the Routing Information Protocol (RIP) and gated.

- **External Gateway Protocol (EGP)**

Using EGP, routers in one autonomous system can advertise the internet addresses of networks in the system to routers in other autonomous systems. Essentially, EGP allows internets to communicate with the Internet backbone.

Routing Information Protocol

Routing Information Protocol (RIP) is an IGP that specifies how routing information passes between routers in an autonomous system. The program that implements RIP is **routed**.

Using RIP, a router periodically broadcasts its routing tables to adjacent routers. The routing tables include information on destination networks and the number of hops required to reach the destination.

An internet **hop** is a measure of distance between two points on an internet. The **hop count** refers to the number of networks a datagram travels on its way from the source host to the destination host. If the destination host for an IP datagram is on the same network as the source host, the hop count is one, because the datagram travels on only one network. If the destination host is on another network, the hop count is two, because the datagram travels from one network, through a router, to a second network.

The number of routers between the source and destination hosts is equal to the hop count minus one. Therefore, if the hop count is 3, the number of routers between the source and destination hosts is 2.

Internet Control Message Protocol (ICMP)

When an internet and all its components are functioning well, routing occurs without errors. However, when software or hardware problems exist, a router may not be able to route or deliver messages, and the network needs a way of reporting errors so that the problems can be resolved. **Internet Control Message Protocol ICMP**, a required part of IP, performs this function. In addition, ICMP provides route change information from IP routers to hosts. The route change information contains alternate routes to use when hosts on the route are unavailable.

ICMP messages occupy the data portion of IP datagrams. However, the ICMP part of the datagram is exchanged only between the internet software on the sending host and the internet software on the receiving host. The user's application never encounters ICMP messages directly.

Overview of Data Communications Networks

What Is an Internet Network?

ICMP messages are organized into types, which cover a wide range of routing problems or failures, including the following:

- Requests to confirm reachability of hosts (echo requests using ping)
- Unreachable destination
- Congestion control
- Route changes
- Forwarding timeouts
- Time information requests
- Requests for information such as network addresses and subaddress masks

Transport Protocols

Transport protocols address process-to-process communications. This section discusses two transport protocols: Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

Transmission Control Protocol

The four-layer Internet Protocol model in Figure 1–17 shows the **Transmission Control Protocol (TCP)** as the layer between IP and application layers. TCP is not software, but rather a communication protocol—a standard according to which communication software is implemented.

TCP's primary function is to provide reliable process-to-process communications. To do this, TCP passes data from IP to users' applications, and from users' applications to IP. TCP provides full-duplex (two-way) connections so that hosts can exchange large volumes of data efficiently.

TCP's reliable stream delivery service is especially important for application programs, which often need to send large amounts of data. Without TCP, each application program would have to include error detection and recovery to ensure that data is properly sent.

In providing reliable process-to-process communications, TCP does the following:

- Transfers data

TCP transfers streams of octets in both directions, and packages octets into segments for IP.
- Ensures reliability by performing the following functions:
 - Recovering from damaged, lost, duplicated, and out-of-order octets by acknowledging data packets
 - Retransmitting data when it does not receive acknowledgment within a certain time
 - Using checksums to ensure that the data has not been damaged
 - Checking the sequence number of each octet, and discarding any duplicate octets
 - Using sequence numbers to order the data properly
- Provides flow control

Flow control allows the receiver to govern the amount of data the sender sends. When the receiver gets data, it acknowledges the data received, and specifies the next range of acceptable sequence numbers the sender can send.

Overview of Data Communications Networks

What Is an Internet Network?

- Provides multiplexing
Multiplexing allows multiple processes on a single host to use TCP simultaneously.
- Creates, maintains, and terminates connections between processes on separate hosts

User Datagram Protocol

User Datagram Protocol (UDP) is the internet standard protocol that allows an application program on one host to send a datagram to an application program on another host. UDP provides unreliable, connectionless delivery service using IP to transport messages among hosts.

UDP is similar to TCP in that it provides a mechanism for user applications to communicate with IP. UDP differs from TCP in that UDP is a very simple protocol, and is entirely dependent upon IP's best effort to provide any reliability. UDP does not guarantee delivery, occasionally generates duplicate data packets, and may send data in the incorrect order. However, layers above UDP can create reliable services using UDP.

The difference conceptually between UDP and IP is that unlike IP, UDP messages include a protocol port number, allowing the sender to distinguish among multiple destinations (application programs) on the remote host.

A common use of UDP is in the Network File System (NFS).

Other Protocols

This section discusses the Address Resolution Protocol (ARP), which enables internet hosts to determine the physical addresses of other hosts on the same network. It also discusses the Simple Network Management Protocol (SNMP), used to manage TCP/IP networks.

Address Resolution Protocol

As stated in Section 1.3.4, internet addresses identify network connections. They do not identify the host itself, nor do they identify the physical network address. To send datagrams to another host on the same network, the sending host needs to know the receiving host's physical address. The **Address Resolution Protocol (ARP)** is the means by which hosts determine the physical network addresses of other hosts on the same network.

When a host wants to communicate with another host on the same physical network, the sending host broadcasts an ARP message requesting the corresponding physical address for the destination host's internet address. Although every host on the same physical network receives this request, only the host whose physical address has been requested responds—thus binding a physical address with the internet address.

The ARP protocol helps minimize the expense of using broadcast messages to obtain address mapping information by providing storage or **cache** for the bindings most recently requested by a host. Because the bindings are stored locally, the number of broadcast messages sent around the network for ARP requests are minimized.

Overview of Data Communications Networks

What Is an Internet Network?

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) manages TCP/IP networks, and is currently an interim standard for network management in TCP/IP networks.

SNMP specifies a protocol that allows you to monitor and control TCP/IP networks through the use of the following network management entities:

- Network management stations
- Network elements
- Management agents

Network management stations execute management applications, which monitor and control network elements. **Network elements** are devices such as hosts, routers, and terminal servers. Each network element uses a **management agent** to perform the network management functions requested by the network management stations. SNMP communicates management information between the network management stations and the agents in the network elements

For more information on SNMP, see RFC 1098, *A Simple Network Management Protocol (SNMP)*.

Network Troubleshooting Methodology

This chapter presents a methodology for solving network problems thoroughly and efficiently using a structured problem-solving approach. Any attempt at problem solving benefits from the use of a structured approach, and the increasing complexity of computer networks makes structured problem solving essential.

As you become adept at network problem solving, you may be able to streamline this process by combining some of the steps. However, for newcomers to network troubleshooting, or for new network problems, the structured approach helps get you started in the right direction.

Knowing Your Network

During network troubleshooting, it helps to have a good understanding of how your network operates under normal circumstances. Understanding normal operation helps you recognize abnormalities in performance or operation that signal potential problems.

This section explains what you need to know about your network, including the topology, architectures, normal performance, and normal use.

Topology

Section 1.1.2 defined network topology as the physical and logical location of components in a network. In network troubleshooting, up-to-date maps of the physical and logical locations of all the devices on the network are critical.

Maps can help you understand the extent or overall impact of a failure, and can help you isolate problems to a particular LAN, LAN segment, or even a service on the segment.

You can do basic tests from your office to isolate the source of a problem, and then match your findings to the problem device on the network map. An accurate network map allows you to go directly to the correct physical source of the problem and begin solving the problem. Without an accurate network map, you can waste valuable time looking for the problem device rather than solving the problem.

The following software tools can help you maintain topological information and network maps:

- DEC Extended LAN Management Software (DECelms)
- NMCC/VAX ETHERnim

Network Troubleshooting Methodology

Knowing Your Network

- NMCC/DECnet Monitor
- DECMcc Management Station for ULTRIX

If these tools are not available, generate network maps manually by walking around your site and recording the location of each device. Regardless of the method you use to generate the maps, the maps must be up-to-date to be of value. Be sure to monitor any changes in the circuits or devices of your network, and to regularly update your maps with the changed information.

Architecture

You need to know the architectures used on your network, the distinctions between the architectures, and the functions of the various architectural layers. Refer to Chapter 1 for information on DECnet, TCP/IP, and Ethernet architectures.

Using DECnet Architectural Information in Problem Solving

Understanding both the architectures used on your network, and the functions of each layer of an architecture enables you to identify and isolate problems quickly and easily. This section describes two typical DECnet problems and explains how understanding architectural information can help isolate the source of the problems.

Recall from Section 1.2.5, the physical and data link layers of DECnet are responsible for creating an error-free path across a communications link between adjacent systems. The routing and end communications layers ensure that data is delivered undamaged and in the right sequence to the correct destination system. Above these layers, appropriate protocols are responsible for interpreting the data that arrives from a sending node.

The following examples show how to use this information to solve "Remote node is not currently reachable," and "Login information invalid at remote node."

- "Remote node is not currently reachable"

When a remote node is unreachable, the lower layers of the architecture (physical, data link, and routing) are probably involved, because you cannot establish a connection to the remote node. You can immediately focus on problems involving these layers.

First, verify that the remote node's name and address is correctly defined on the local node, and that the user's local node is on the network. Then, you can verify that the problem is in the lower layers of the architecture by trying to establish a connection to another node on the same LAN as the unreachable node. If you cannot reach either the original or the second node, the problem is probably a LAN segmentation problem, which involves the physical or data link layers. If you can reach the second node, but not the original remote node, the problem is probably still at the physical or data link level. However, in this case, it is related to the original remote node rather than the LAN.

- "Login information invalid at remote node"

With the message "Login information invalid at remote node," you can rule out problems with the physical and data link layers, because this message comes from the remote system, and the lower-level connectivity is intact. The problem is in the session control layer, where application dialogs start.

Specifically, "Login information invalid at remote node" is a DECnet error in the session control layer. The error is in DECnet because remote login is a DECnet operation. The session control layer of DECnet is involved because the access control information is invalid and the system-dependent portion of the logical link creation failed. This means that the initial connection was made, but the remote node could not complete the connection because of problems with the access control information.

After you identify the layer at which the problem occurred, you can select the correct tools to help solve the problem. In the case of "Login information invalid at remote node," because the problem involves process creation or access control for accounts and proxies, you use the Authorize and NCP utilities to resolve it.

Using TCP/IP Architectural Information in Problem Solving

This section describes how understanding architectural information can help you solve two TCP/IP problems: "Network is unreachable," and "Login incorrect." For a review of the four-layer internet protocol model, refer to Section 1.3.2

- "Network is unreachable"

The "Network is unreachable" problem occurs at the internet layer of the internet protocol model. The internet layer routes messages based on the network portion of the internet address. The "Network is unreachable" error occurs because some problem prevents the internet layer from routing the message to the destination network. When an IP router receives a datagram that is destined for an unreachable network or host, the router notifies the source host using an ICMP message.

To troubleshoot the "Network is unreachable" problem, use the traceroute utility or the netstat -r command as described in Chapter 3 to determine the path of communications through the network, and locate the problem area. You can also get path information using the network ID portion of an internet address and manually checking the routing tables from point to point until the error occurs.

Network Troubleshooting Methodology

Knowing Your Network

- "Login incorrect"

With the "Login incorrect" problem, the internet layer has performed its task of routing the message to the destination, and that the TCP or UDP layer has created a link to the remote host. Thus, the problem lies with the application layer.

To troubleshoot the "Login incorrect" problem, check the login information that the user supplied, against the information contained in the remote node's `/etc/passwd` file. Correct the information in the `/etc/passwd` file if necessary.

Performance

If you understand the normal ranges of performance in your network, you can better determine when the network is performing poorly due to a network problem. You can use tools such as the LAN Traffic Monitor for LANs, and NMCC/DECnet Monitor for WANs to accumulate information on historical performance and error characteristics. You can also use NMCC/VAX ETHERnim and the Network Control Program (NCP) to accumulate performance information. For TCP/IP networks, you can use DECMCC Management Station for ULTRIX to gather performance information.

In understanding how network performance relates to network troubleshooting, the important concepts are thresholds and usage peaks. **Threshold** is the maximum value set for a parameter. **Usage peak** is the maximum level of activity the network can withstand before performance is adversely affected.

When thresholds are reached or when activity reaches a peak, transient and intermittent performance problems may begin to surface. Performance problems can be among the most difficult to isolate, primarily because they are often transient and intermittent. (See Section 2.6.2 for more information about transient and intermittent errors.)

For example, when traffic and queuing requests increase, bandwidth thresholds can be reached and performance problems can occur. NMCC/DECnet Monitor's graphics display can help in uncovering problems related to bandwidth and utilization thresholds, and can help identify whether insufficient bandwidth is the cause of performance problems.

Maintaining historical performance data is essential for troubleshooting transient and intermittent problems. For example, if you have a transient problem, you can compare the historical performance data to the current performance data, and evaluate the differences to help isolate the source of the problem. Historical performance data is also very useful in trend analysis for network growth and network planning.

Typical Use

To understand the typical use of your network, you need to be aware of the following:

- Applications running on your network
- Times of day when those applications are running
- Peak performance periods

You need to know the types of applications used on your network, and the typical traffic patterns that the applications produce. For example, some applications may cause performance behavior that is different from normal operation, but still within accepted bounds for that application. Understanding this helps you determine when a change in performance is acceptable, and when the change is not acceptable and requires action on your part.

Also be aware of any unusual performance that a particular application causes, and the possible effects it may have on other applications running at the same time. Understanding the performance of the applications in your network allows you to predict peak usage times, anticipate problems before they occur, and design your network to accommodate the needs of users who require these applications.

Overview of the Network Troubleshooting Methodology

Users expect reliable service and fast resolution to network problems. Your role as a network troubleshooter is to ensure high network availability—in other words, to minimize the amount of time the network resources are down. This section shows you how to apply the knowledge of your network in a methodical way to resolve any network problem.

The network troubleshooting methodology consists of the following steps:

1. Obtain and refine a problem statement.
2. Gather information about the problem.
3. Analyze, interpret, and classify the information.
4. Solve the problem or refer to the responsible person.
5. Verify the solution.
6. Clean up any parameters or files created for testing, and remove any test equipment such as loopback connectors.
7. Verify the solution again.
8. Document the problem and archive the solution.

Figure 2–1 shows the relationships among the steps of the network troubleshooting methodology. The flowchart is based on the assumption that you understand your network's topology, architectures, software, performance, and normal use, and can apply this knowledge throughout the process.

The remaining sections of this chapter describe each step of the troubleshooting methodology in detail.

Network Troubleshooting Methodology

Overview of the Network Troubleshooting Methodology

Figure 2–1 Network Troubleshooting Methodology

Applying The Methodology to Network Problems

Gathering, analyzing, interpreting, and classifying information are steps that you may need to perform repeatedly throughout the troubleshooting process. The following example shows how these steps apply to solving three similar problems on DECnet and TCP/IP networks:

- "Remote node is not currently reachable." (DECnet)
- "Connection timed out." (TCP/IP)
- "Host unreachable." (TCP/IP)

Network Troubleshooting Methodology

Overview of the Network Troubleshooting Methodology

The example uses the network configuration shown in Figure 2–2, where computer Z cannot reach computer A.

Figure 2–2 Unreachable DECnet Node or TCP/IP Host

1. For the DECnet problem, "Remote node is not currently reachable," ensure that A's address is correctly defined in Z's node database, and that node Z is on the network.

(No corresponding step for the TCP/IP problems.)

2. For the DECnet problem, use the following command to verify that Z cannot reach A:

```
$ DIRECTORY A::
```

If this command fails, you receive the error, "Remote node is not currently reachable." Because an unreachable node is a problem in the physical, data link, or routing layers of DNA, you can focus on problems in those layers, and rule out the user's application as the source of the problem.

This step helps you refine the problem statement so that you can focus your troubleshooting efforts appropriately.

For the TCP/IP problems, use the ping command from host Z to determine whether host A is reachable.

Host A is connected to the same physical network as host Z. Therefore, if host A does not return the message "hostname is alive," then the problem is in the lower layers of the IP model. The lower layers are responsible for routing and for transmitting data across the network media.

3. Assuming that the LAT protocol is active, verify whether or not the problem is specific to the LAN by trying to connect to the node through the LAT.

If you can reach the node through the LAT but not through DECnet or TCP/IP, you can focus your efforts on the lower layers of DECnet or TCP/IP on node or host A.

In this step, you gather information, then analyze, interpret, and classify the information to help determine the cause of the problem, and narrow the focus of your efforts.

Network Troubleshooting Methodology

Overview of the Network Troubleshooting Methodology

4. If you cannot reach the node through the LAT, then use an accurate LAN map, the NMCC/VAX ETHERnim display, or DECmcc Management Station for ULTRIX as a guide to the network, and to try to establish connections to other nodes on the same LAN segment as node A.

If you cannot reach any other nodes on the segment (nodes B,C, and D), the problem is a LAN segment problem. If you can reach some nodes on the segment (for example, node D) but not others (nodes A,B,C), the problem may be related to the cable between the DELNI and the H4000 transceiver that connects the DELNI to the Ethernet.

If all connections to the segment fail, the problem is probably related to a bridge or repeater failure, a cable failure, or a babbling device. If only node A fails, the problem is specific to node A.

In this step, you gather more information, then analyze, interpret, and classify the information to help determine the cause of the problem.

The preceding example illustrates the following points:

- Knowing DNA and TCP/IP helps you identify the problem as a lower-layer problem.
- Knowing LAN architecture helps you identify the problem as LAN-specific.
- Knowing the topology and how to use a network map helps you isolate the problem further.

Network Troubleshooting Methodology

Overview of the Network Troubleshooting Methodology

- Knowing how to use network tools helps you perform tests.
- Testing connectivity to other remote nodes helps you eliminate reachable nodes from the troubleshooting effort, and isolate the problem to a LAN segment.

Detecting a Problem

A network problem is any failure of hardware, software, or communications media that prevents users from accessing network resources. Network problems may present themselves as explicit error messages, unexpected behavior during normal operation, or a lack of response to a normal operating request. You may be alerted to an existing or potential problem by users, through unusual behavior on the network, through error conditions displayed by network management tools such as NMCC/DECnet Monitor, NMCC/VAX ETHERnim, DECMCC Management Station for ULTRIX, or through your normal monitoring of network operations.

Once you are alerted to the problem, the troubleshooting process begins.

Obtaining and Refining the Problem Statement

The first step in the troubleshooting process is to ensure that you have a clear, unambiguous understanding of the problem. A clear understanding helps you investigate the problem correctly the *first* time.

When you first encounter a problem, ask the following preliminary questions of yourself or the user who reported the problem:

- Is the problem accompanied by one or more error messages?
If so, what are the messages, and in what order are they displayed?
- Can you re-create the error?
- What system or systems are involved?
- What operation were you performing when the error occurred?
- What application or network object is involved?
- At what time did the error occur?

Use the answers to these questions to formulate a clear problem statement.

Note

Any network problem could be caused by a user error, such as mistyping access information. However, the network troubleshooting process does not specifically address user errors. Be sure to confirm that the user did not make an error before you start troubleshooting the problem.

Gathering Information

After you have a clear problem statement, gathering additional information helps you narrow both the scope of the problem and the possible causes. As shown in Figure 2-1, you may need to gather information several times throughout the process to properly isolate the source of the problem.

Use the following questions to guide you through the information gathering step:

- **Is NMCC/VAX ETHERnim generating data on the problem?**

Network Troubleshooting Methodology

Gathering Information

If the background tasks are set up for polling, look for nodes displayed as dots and dashes. This condition might indicate a topology problem, such as a problem with a repeater or DELNI.

- **If yours is a DECnet network, is NMCC/DECnet Monitor displaying alarms?**

If the map displays red conditions, then error rates or utilization might be too high. These conditions tend to slow performance and cause links to drop, thus creating reachability problems. NMCC/DECnet Monitor indicates the specific circuit, wire, or router involved in the problem.

- **Does the LAN Traffic Monitor data show anything unusual or significant, such as a very high Top Talker?**

A very high top talker could indicate one of two conditions:

- If several links (particularly LAT and Local Area VAXcluster links) are dropping, then the top talker may indicate a babbling device. (See Chapter 5 for more information on babbling devices.)
- If the protocol involved is one that sends out a substantial amount of broadcast traffic, then top talker might indicate a broadcast storm. (See Chapter 5 for more information on broadcast storms.)

- **Is the failure a hard, inconsistent, intermittent, or transient error, as defined in Section 2.6.2?**

See Section 2.6.2 for information on these types of errors.

- **What protocols are on the network?**

This information is particularly important on Ethernet networks that run multiple protocols.

For example, on Ethernet networks running DECnet, DECnet must be the first protocol that starts. This is because DECnet changes the hardware address to a hexadecimal equivalent of the DECnet address when it starts. If other protocols on the controller start before DECnet, the other protocols use the hardware address, preventing DECnet from changing the address as required and causing an error. (See "Invalid Parameter Value" in Chapter 5.)

On internet networks, problems with the Address Resolution Protocol (ARP) can cause network errors. For example, faulty ARP tables can cause problems with direct routing on a network.

- **In what architectural layer is the problem?**

By determining the architectural layer involved, you can often determine the extent of the problem (node or host, LAN, WAN). For example, DECnet session layer problems are always node problems on either the local or remote node.

- **What applications are affected?**

You need to know the applications affected so that you can replicate the problem. If you know that an application has certain requirements, you can check to make that sure those requirements are met. For example, the application may require proxy access. If proxy access is not properly set up, the application may not work.

In the case of DECnet applications, you may notice that some applications work, and others do not. To begin to solve this problem, start investigating the session layer, because this is where the differences in applications are most apparent. Remember that protocols such as LAT are more time sensitive than DECnet. As a result, LAT connections to applications may drop before other connections.

After you solve the problem, use the application to confirm that the solution you used to fix the problem worked.

- **Can you determine the extent of the problem as a node, host, LAN, or WAN problem as described in Section 2.6.1?**

If yes, see the appropriate area in Section 2.7.

If no, continue with the next question, and assume the problem is in the network software. Use TCP/IP or DECnet (depending on the the software involved in the problem) to troubleshoot the problem. You assume the problem is in DECnet or TCP/IP, because network software problems can occur at the node, host, LAN, or WAN level. Other protocols, such as LAT, and bridge are LAN-specific.

- **Is the error limited to a group of nodes or hosts on a LAN segment?**

If yes, then it is a LAN problem. See Section 2.7.2.

If no, then continue with the next question.

- **Does the error occur when executed from the node or host in question to itself?**

For example, if it is a DECnet problem, does the error occur if you try to use the SET HOST command from the node in question to itself? Does the error occur if you use the command DIRECTORY *node-id::* to display the directory of the default DECnet account?

If it is a TCP/IP problem, does the error occur if you try to use the rlogin command to make a network connection from the host in question to itself?

If yes, then the problem is node-specific or host-specific. See Section 2.7.1.

If no, then treat the problem as a WAN problem until you can isolate it to either a WAN, node-specific, or host-specific problem using routing path trace procedures. See Section 2.7.3.

Analyzing, Interpreting, and Classifying Information

This step of the troubleshooting process involves methodically evaluating the information you collect. As shown in Figure 2-1, you may need to analyze, interpret, and classify information repeatedly as you gather more specific information about the problem you are solving. The continual information gathering throughout the troubleshooting process, followed by analysis, interpretation, and classification of the information helps you isolate the source of the problem and allows you to eventually solve the problem.

During this step, keep in mind the goal of isolating the problem physically to the node, host, LAN, or WAN levels, and logically to a specific architecture and layer of that architecture.

In *analyzing* the information, you examine the elements of information and the relationship of each element to the others.

In *interpreting* the information, you evaluate the elements of information and the relationships among the elements to get an understanding of the problem.

Network Troubleshooting Methodology

Analyzing, Interpreting, and Classifying Information

In *classifying* the data, you group the information so that you can rule out certain problems and begin to isolate the source of the problem.

This section discusses classifications of errors by the following categories:

- Extent of the problem
- Types of network errors
- Sources of network errors

Extent of the Problem

The extent of a network problem refers to the bounds of its disturbance on the network. For example, some problems interfere with correct operation of a single node or host, while other problems interfere with correct operation of multiple nodes or hosts, an entire LAN, or even an entire WAN.

Classifying the extent of the problem continues to focus your efforts, limits the scope of your investigation, and helps you select an appropriate approach to solving the problem. After you classify a problem as specific to a single node or host, a LAN, or a WAN, you further define and classify the problem within that area.

The following sections define node or host, LAN, and WAN problems.

Node or Host Problems

A node or host problem is limited to a specific node or host.

A VMS node problem generally involves parameters set in AUTHORIZE, NCP, LATCP, and SYSGEN.

A host problem can involve incorrect information in tables or files intended to provide routing information, such as the following:

- Routing tables
- /etc/hosts file
- /etc/hosts.equiv file
- local.hosts file
- .rhosts file
- BIND name server host files
- Yellow Pages database files

In addition, host problems may involve incorrect parameters set through the ifconfig program.

LAN Problems

A LAN problem occurs only in an Ethernet environment. A LAN problem is specific to the following:

- Ethernet protocols, such as ARP, LAT, and SCA
- Ethernet design rules, such as repeater rules, bridge rules, and cabling rules (see Appendix A)
- Ethernet hardware, such as bridges, repeaters, connectors, and cabling

Network Troubleshooting Methodology Analyzing, Interpreting, and Classifying Information

As you begin to isolate the problem source, you may find that a LAN problem is, in fact, related to a specific node. Nevertheless, it is helpful to begin classifying the extent of the problem as LAN-wide until you have more information to determine otherwise.

WAN Problems

A WAN problem extends beyond the Ethernet into a point-to-point DDCMP environment. For example, DECnet-related WAN problems include circuit down, router unavailable, partitioned area, and remote node access problems. WAN problems on TCP/IP networks include down or malfunctioning IP routers, problems with links between routers, routing protocol problems, and name server problems.

As with LAN problems, you may find that a WAN problem is, in fact, related to a specific node. Again, it is helpful to at least begin classifying the extent of the problem as WAN-wide until you have more information to determine otherwise.

Types of Errors

The following sections define four types of errors:

- Hard
- Inconsistent
- Intermittent
- Transient

Hard Errors

Hard errors are consistently reproducible, and consistently produce the same error message or symptom when reproduced under the same circumstances.

Inconsistent Errors

When several different error messages or symptoms ultimately have the same underlying cause, the errors are considered inconsistent. Inconsistent errors generally involve several protocols or several layers of an architecture. The different error messages result from the ways different applications encounter the problem in the protocols and the architectural layers. You can usually reproduce an inconsistent error.

For example, on a DECnet network, when trying to mail a file to a remote node, a user may get the following error message:

```
%MAIL-E-SENDERR, error sending to user SMITH at NODEID  
%MAIL-E-PROTOCOL, network protocol error  
-SYSTEM-F-LINKABORT, network partner aborted logical link
```

However, if the user tries to copy the file to the remote node, the user receives the following error:

```
%COPY-E-OPENOUT, error opening NODEID::USER$31:[SMITH]MYFILE.TEMP;1 as output  
-RMS-F-FUL, device full (insufficient space for allocation  
%COPY-W-NOTCOPIED, USER$25:[JONES]MYFILE.TEMP;1 not copied
```

Both of these problems are related to a lack of disk space for the default DECnet account. However, the error messages are different because the means of accessing the remote node are different—in one case through mail, and in the other through remote file access.

Network Troubleshooting Methodology

Analyzing, Interpreting, and Classifying Information

Intermittent Errors

An intermittent error is one that shows up occasionally, and that always displays the same error message or symptoms for the same circumstances. You can occasionally reproduce intermittent errors.

For example, setting up a cluster alias on a DECnet node improperly may result in intermittent errors. An error occurs when an improperly configured node in the cluster is requested to perform an alias node function. The user receives an error only when the request to the cluster goes to the improperly configured node. For example, if a remote user tries to establish a connection to the cluster, the connection may fail if the node receiving the request is the misconfigured node. However, if the user tries the request again, and the node receiving the request is another, properly configured node, the request succeeds.

Intermittent errors may also result when threshold values for various parameters are reached. These thresholds are usually sufficient for normal use, but during peak use, the thresholds may be reached and errors can result.

Transient Errors

Transient errors are those that occur only occasionally, and can rarely be reproduced. Because you cannot reliably reproduce transient errors, they are by far the most difficult errors to troubleshoot.

As with intermittent errors, transient errors may result when threshold values for various parameters are reached. Because transient errors tend to occur at peak usage times, historical performance data is helpful in determining the cause of the problem.

Sources of Errors

The following sections describe potential sources of network problems, including the following:

- User errors
- Hardware errors
- Software errors
- Configuration errors
- Interoperability errors

Understanding the possible sources of errors enables you to ask questions during the troubleshooting process that help to narrow the scope of the problem, and allows you to quickly isolate the source of the problem.

User Errors

User errors include typing errors, command syntax errors, improper use of hardware or software due to an unclear understanding of its function, and poor applications programming. User errors can result in hard, inconsistent, intermittent, and transient failures.

Hardware Errors

Hardware errors include failed devices, loose connections, faulty or noisy circuitry, and lack of power. Most of these errors result in hard errors, although loose connections and noise on the line can cause intermittent problems. Noisy or dirty circuitry may also cause transient errors, only occurring at certain traffic levels or with certain bit patterns.

Network Troubleshooting Methodology

Analyzing, Interpreting, and Classifying Information

Software Errors

Software errors involve improperly configured software, thresholds being reached, or limitations of the product, such as use of the product in ways other than those intended and specified. Misconfigured software tends to result in hard errors. Thresholds being reached and product limitations tend to result in intermittent and transient errors.

Configuration Errors

Configuration errors on a LAN result from failing to conform to recommendations for product installation and use. For example, errors may result from failing to conform to installation and usage guidelines for Ethernet. Performance problems may result from failing to conform to recommendations for logical network configurations.

On internet networks, improper definition of the network mask or broadcast address can cause communication problems for the local host, and may result in problems, such as broadcast storms, that affect the rest of the network.

Configuration errors can result in intermittent problems or hard errors.

Interoperability Errors

An interoperability error indicates a problem resulting from incompatibilities between protocols. Interoperability problems can occur as the result of either of the following:

- Incompatibilities between different implementations of the same protocol
This is the most common type of interoperability problem. It results from the way different vendors interpret and implement the same protocol specification. If the vendors vary in their interpretations of the specifications, their two implementations may be unable to communicate with each other, or may communicate inefficiently.
- Incompatibilities between entirely different protocols
This type of interoperability problem occurs when one protocol causes problems with other protocols on the same network. Multiple protocols can successfully use the same media, however, occasionally one protocol adversely affects other protocols using that media. This type of interoperability problem usually occurs on networks that have both multiprotocol and broadcast capabilities. For example, a broadcast storm (See Chapter 5) can affect DECnet nodes even though DECnet nodes do not use the broadcast address.

Isolating the Source of the Problem

Quick and efficient fault isolation is key to resolving network problems. **Fault isolation** is the process used to determine the source of a problem. In many cases, fault isolation is the most difficult and time consuming part of resolving a network problem. The more orderly and efficient the fault isolation process, the quicker you can return normal network service to your users.

In some cases, isolating the source of the problem may be all that you can do to resolve a problem. For example, you may determine that the problem lies with a common carrier, or requires hardware repair, or involves system parameters on a node to which you do not have access. In these cases, you must turn over responsibility for resolving the problem to the appropriate individual. However, by properly isolating the problem first, you allow other individuals to focus on the repair that needs to be done, and help to minimize the overall downtime of network resources.

Network Troubleshooting Methodology

Isolating the Source of the Problem

In isolating the source of a problem, you use the information you accumulated to limit the scope of your investigation from the broadest possible categorization of the problem to the narrowest—eventually focusing on the specific cause of the problem.

This section discusses tools and methods to use for isolating problems to the node, LAN, and WAN levels. Figure 2–3 shows the problem isolation step in flowchart form.

Isolating to the Node or Host Level

Most problems, including LAN and WAN problems, can eventually be tracked to problems on a specific node or host. To isolate suspected node problems, retry the failed operation to other nodes to determine if the problem affects only one node or multiple nodes. For example, if you can get from node A to node C, but not to node B from nodes A or C, then the problem is most likely in node B.

After you isolate a problem to a specific node, you need to determine what exactly on that node is causing the problem. Usually, the problem involves DECnet, TCP/IP, MOP, or LAT protocols, or faulty hardware.

When you troubleshoot node problems, keep in mind the protocols and the architectural layers involved in the problem. When you understand which protocol or architectural layer is involved, you automatically narrow down the potential factors contributing to the problem, and can focus on the true cause.

Likewise, error messages allow you to refine your understanding of a problem. With practice, you can use error messages to quickly eliminate potential sources of the problem, and focus on a specific problem area. Examples for DECnet and TCP/IP networks follow.

DECnet Networks

Once you know that the OPCOM message, "Circuit on-starting" indicates a hardware, cabling, modem, or circuit problem, you can use the error log file (see VMS Error Log File) to gather more information about the potential hardware problem causing it. Loopback tests (see NCP Loopback Tests) provide additional information that can help you isolate the problem to the cabling, modem, or circuit.

The error message, "Device not mounted," indicates that DECnet has not yet started. Errors that occur while starting DECnet generally result from the Ethernet device being unable to accept a DECnet address because another protocol is already using the hardware address. To resolve this problem, you need to stop the other protocols, make sure DECnet starts first, then restart the other protocols.

The error messages, "Login information invalid," "Network object unknown," and "Network partner exited," are all DECnet software errors that occur in the upper layers of DECnet. Knowing this allows you to eliminate potential approaches to the problem and to focus on using the tools most likely to help you isolate the source of the problem for these errors. In resolving these types of problems, you may need to run AUTHORIZE to access the SYSUAF file, or NCP to access the network object database. You may also need to use process creation procedures.

Figure 2–3 Isolating the Source of the Problem

TCP/IP Networks

For TCP/IP hosts, the error messages "Login incorrect," "Permission denied," and "Connection refused" all point to a host problem on either the local or remote host, not a problem between the hosts. In each of these cases, the remote host refuses access, indicating that communication with the remote host has been established.

Likewise, the error message, "Host unknown" points to a host problem. However, in the case of "Host Unknown," the message implies there is a problem with the name service, or that the user is attempting to use an incorrect host name.

The TCP/IP tools that you can use to solve these problems include ifconfig and netstat.

Table 2–1 summarizes the most effective tools for isolating node or host problems.

Table 2–1 Tools for Node or Host Problems

Type of Problem	Tools to Use
DECnet and MOP	NCP, AUTHORIZE, SYSGEN, AUTOGEN, ETHERnim, DECelms, netserver log files
TCP/IP	ifconfig, netstat command, ping command, traceroute command
LAT	LATCP, SYSGEN, AUTOGEN
Hardware	Error log files, loopback connectors

Network Troubleshooting Methodology

Isolating the Source of the Problem

Isolating to the LAN Level

Because of the nature of bus architectures such as Ethernet, it can be difficult to isolate a problem device in the event of a failure. However, tools such as NMCC /VAX ETHERnim, combined with a good knowledge of your network topology, can be very helpful in locating the problem source.

After you isolate a problem to a LAN, you can further isolate the problem within the LAN by determining which LAN segment has the problem. When you know the LAN segment involved in the problem, you can focus your efforts on determining which node or nodes on that segment are causing the problem or are being affected by the problem. This is where your knowledge of the topology comes into play. You need to know the segments that make up your LAN, and the individual nodes on each segment.

Most LAN problems result from problems with a single device. For example, a LAN babbler is one node transmitting incorrectly, and a LAN segment down is a bridge or repeater failure or a short-circuited connection.

Other important information to keep in mind for solving LAN problems includes Ethernet cabling rules (see Appendix A for guidelines on Ethernet cabling), Ethernet-specific protocols, and Ethernet concepts (CSMA/CD, slot time, and round trip propagation delay). This information helps you understand why a problem has occurred, and how to prevent future problems from occurring.

For example, LAN performance problems may result from configuration errors. Using your knowledge of Ethernet concepts, cabling, and protocols, you can configure your network properly by locating devices on the LAN for optimal performance.

LAN problems can also involve LAT or downline loading protocols such as MOP, and may be caused by problems in a specific area or with a group of nodes. LAN problems can include segmentation problems, cabling problems, or a terminal server problem.

Table 2-2 summarizes the most effective tools for isolating LAN problems.

Table 2-2 Tools for LAN Problems

Type of Problem	Tools to Use
LAT	LATCP, TSM, DECserver commands
MOP	DECelms, OPCOM, ETHERnim, TSM, NCP
Cabling	DECelms, ETHERnim
TCP/IP	arp command, netstat command, ping command
Node or host	See Table 2-1

Isolating to the WAN Level

WAN problems only occur in point-to-point environments, and all WAN problems involve circuit, line, or router problems.

The first thing to do when you suspect a WAN problem is to try the failed operation to other nodes on your LAN to eliminate the possibility that the problem is actually LAN-based rather than WAN-based. If you can complete the operation on the LAN, try the operation to other nodes outside the LAN, starting with nodes just outside (adjacent) to the LAN, and working your way out to nodes

Network Troubleshooting Methodology Isolating the Source of the Problem

further and further from the LAN. This helps you determine the extent of the failure.

If you find you cannot complete the operation to other remote nodes, the next step is to trace the routing path to isolate potential causes of the problem. (See Routing Path Trace Procedures for information on how to trace a routing path.) NMCC/DECnet Monitor also helps in quickly pointing out potential WAN problem areas by highlighting failed routers or circuits on its display.

For circuit problems, loopback tests (see NCP Loopback Tests) help to determine which components in the path are functional, and which components in the path are not functional.

Router problems can cause unreachable node problems, area partitioning, and circuit state problems. Router problems are mostly node problems, because a router is a node designated to process DECnet traffic from one point to another.

For TCP/IP problems such as "Connection timed out" and "Network is unreachable," you can manually trace the route using netstat or you can use the traceroute utility to isolate the source of the problem.

Using the traceroute utility and your knowledge of the network topology, you can determine where connectivity has been lost. However, in some cases, IP routers may be invisible to the traceroute utility. If this is the case, you can use the ping command to manually check if each of the routers along the path to the destination is running and reachable on its incoming link.

Usually the problem involves either an incorrect or missing entry in the routing tables, or a circuit problem. As with any problem involving remote systems, access control restrictions may require that you work with other network and system managers to resolve these problems.

Table 2-3 summarizes the most effective tools for isolating WAN problems.

Table 2-3 Tools for WAN Problems

Type of Problem	Tools to Use
Circuit	NCP, DECnet Monitor, loopback tests, path traces
Line	NCP, DECnet Monitor, loopback tests, path traces
TCP/IP	ping command, SNMP tools, syslog daemon, traceroute command
Node or host	See Table 2-1
LAN	See Table 2-2

Solving the Problem

At this point, you have formulated a problem statement, gathered and evaluated information, and have isolated the problem sufficiently enough to apply a solution.

If you cannot solve a network problem quickly and network availability is critical, you may find it necessary to implement temporary solutions that enable the network to remain running, while still allowing you to work on the problem. These temporary measures are not ideal solutions, and may cause a decrease in network performance. However, for most users, decreased service during the repair phase is preferable to no service at all.

Network Troubleshooting Methodology

Solving the Problem

For example, assume your 9.6 kbs circuit is down, and the common carrier service will not be able to fix it for at least 48 hours. You could use the 2.4 kbs dialup lines to reestablish network connectivity, although at 25% of normal performance. For example, with such a drop in performance, a file transfer that normally takes one hour, could take as much as four hours or more. Ordinarily, this level of service is unacceptable. However, this decrease may be acceptable on a temporary basis, while the common carrier resolves the problem.

If the required solution is beyond the bounds of your authority or expertise, you may need to seek help from another person.

Verifying the Solution

The purpose of this step is to test the solution to confirm that the problem is correctly and adequately solved. If the solution does not solve the problem, you need to gather more information, and analyze and interpret the data again.

Cleaning Up

After you solve a network problem, and verify the solution, reset and remove all test facilities, and any solutions that have been temporarily installed or enabled for troubleshooting. For example, remove loopback connectors, cancel any requests you might have made for service from outside vendors, and delete test files. This step ensures that any parameters or facilities you used in solving the problem do not interfere with the operation of the network.

Verifying the Solution Again

After you clean up any temporary test facilities and solutions you may have required to solve the problem, confirm that your solution is still correct. If the solution does not solve the problem, you need to gather more information, and analyze and interpret the data again.

Documenting the Problem and Solution

Because network problems can be complicated to solve, and because network maintenance is often a team effort, it is helpful to keep a record of the problems that occur on your network and how you solved them. When a problem recurs, you or any other person responsible for maintaining the network can consult the record and move quickly to the solution.

When documenting the problem, be sure to include the symptoms of the problem, when it occurred, the explanation for its occurrence, how you solved it, and the tools you used to solve it. Also include any temporary measures put in place to enable users to continue working while you solved the problem.

Network Management and Troubleshooting Tools

This chapter discusses the tools available for troubleshooting network problems. It lists the tools alphabetically, gives an overview of each tool, and explains how to invoke the tool, display online help, and exit from the tool. This chapter also discusses the privileges and resources required to run the tools, and the kinds of network problems each tool is best suited for solving. For more detailed information, see the product documentation.

Table 3–1 summarizes information about the network tools, including the operating system on which the tool runs, the environment in which it is usually used (VMS node, ULTRIX host, LAN, or WAN), and its primary uses.

Table 3–1 Network Management and Troubleshooting Tools

Tool	Runs On	Network Environment	Uses
arp command	ULTRIX	LAN	Displays and modifies the internet-to-Ethernet address translation tables used by the Address Resolution Protocol (ARP) on ULTRIX systems
Authorize Utility	VMS	Node	Controls access to VMS systems and allocates resources to users
Breakout boxes	Not applicable	WAN	Analyzes the signals sent between two devices, usually between a DCE and DTE
DEC Extended LAN Management Software	VMS	LAN	Monitors and controls LAN bridges
DECmcc Management Station for ULTRIX	ULTRIX	LAN, WAN	Collects, formats, and monitors IP and DECnet network data from SNMP-based agents and DECnet hosts
LAN Traffic Monitor	VMS	LAN	Collects traffic data for any protocol type on an extended Ethernet
LAT Control Program	ULTRIX, VMS	VMS Node, ULTRIX Host, LAN	Configures and controls LAT protocol on VMS and ULTRIX systems

(continued on next page)

Network Management and Troubleshooting Tools

Table 3–1 (Cont.) Network Management and Troubleshooting Tools

Tool	Runs On	Network Environment	Uses
netstat command	ULTRIX	LAN, WAN	Displays the contents of network-related data structures on ULTRIX systems
Network Control Program	ULTRIX, VMS	Node, LAN, WAN	Configures and controls DECnet-VAX networks, and monitors and tests network resources
NMCC/DECnet Monitor	VMS	WAN	Collects, analyzes, and evaluates network data for Phase III and Phase IV DECnet network nodes, and creates and maintains databases of node and link information
NMCC /VAX ETHERnim	VMS	LAN	Gathers information about Ethernet nodes, verifies that nodes are reachable, displays network topology, and monitors Ethernet traffic
ping command	ULTRIX	LAN, WAN	Determines the reachability of hosts on an internet network
Protocol analyzers	Not applicable	LAN, WAN	Analyzes protocol frames
syslog daemon	ULTRIX	Host	Logs system messages on ULTRIX systems
Terminal Server Manager Software	VMS	LAN	Monitors and controls terminal servers in an extended LAN
traceroute command	ULTRIX	LAN, WAN	Displays the route a datagram takes to a network host
uerf command	ULTRIX	Host	Displays system events including error messages about the hardware and software, and information about system status, startup, and diagnostics

arp command

Overview

The arp command in ULTRIX displays and modifies internet-to-Ethernet address translation tables used by the Address Resolution Protocol (ARP).

The arp command is useful in solving direct routing problems resulting from the following circumstances:

- A source host having incorrect Ethernet address information for a destination host
- Two hosts having the same host name

Required Resources & Privileges

You must log in to the superuser account to use the arp command for changing entries in the internet-to-Ethernet address translation tables.

How to Use

Use the following command syntax to translate an internet address to an Ethernet address:

```
# /etc/arp hostname
```

For example, to get the Ethernet address for an internet host called host1, do the following:

```
# /etc/arp host1  
host1 (16.20.32.2) at aa:0:4:0:8f:11 permanent trailers
```

The system response to this command tells you that the Ethernet address for host1 is aa-00-04-00-8f-11.

Recommended Use

Use the arp command when troubleshooting direct routing problems on an Ethernet. "Connection Timed Out" in Chapter 5 explains how to use the arp command in solving direct routing problems.

Authorize Utility

Overview

The Authorize Utility (AUTHORIZE) is a VMS system management tool for controlling access to the system and allocating resources to users. Using AUTHORIZE, you can create new records and modify existing records in the following files:

- System user authorization file (SYSSSYSTEM:SYSUAF.DAT)
- Network proxy authorization file (SYSSSYSTEM:NETPROXY.DAT)
- Rights database file (SYSSSYSTEM:RIGHTSLIST.DAT)

The preceding list shows the default location of these files. Be aware that logical name definitions for these files (especially in cluster environments) may place these files elsewhere.

Required Resources & Privileges

To use AUTHORIZE, you need the following privileges:

- SYSNAM
- SYSPRV

Use AUTHORIZE only if you are a system manager or if you have the permission of the system manager.

How to Use

To run AUTHORIZE, do one of the following:

If the SYSUAF logical name is defined, enter the following command:

```
$ MCR AUTHORIZE
```

Network Management and Troubleshooting Tools

Authorize Utility

If the SYSUAF logical name is *not* defined, set your default directory to SYSSYSTEM (or the directory where the SYSUAF file is located), and enter the following command:

```
$ MCR AUTHORIZE
```

To get online help while using AUTHORIZE, enter the following command:

```
UAF> HELP
```

To exit AUTHORIZE, enter the following command:

```
UAF> EXIT
```

Recommended Use

For network problems that involve user, network object, proxy, privileged, and nonprivileged accounts, you can use AUTHORIZE to do the following:

- Display information about existing accounts
- Create new accounts
- Modify accounts

Example 3-1 shows how you can use AUTHORIZE to display information that can help you solve the problem, "Invalid Login Information at Remote Node." Example 3-1 shows the kind of information AUTHORIZE displays. In this example, the number of login failures **1** indicates a potential problem.

Example 3-1 Authorize Utility Example

```
$ MCR AUTHORIZE
UAF> SHOW DECNET

Username: DECNET                               Owner:  DECNET DEFAULT
Account:  DECNET                               UIC:    [376,376] ([DECNET])
CLI:      DCL                                  Tables: DCLTABLES
Default:  SYS$SPECIFIC:[DECNET]
LGICMD:   LOGIN.COM
Login Flags: NL:
Primary days:  Mon Tue Wed Thu Fri
Secondary days:                               Sat Sun
Primary 00000000001111111112222 Secondary 00000000001111111112222
Day Hours 012345678901234567890123 Day Hours 012345678901234567890123
Network: ##### Full access #####           ##### Full access #####
Batch:  ----- No access -----           ----- No access -----
Local:  ----- No access -----           ----- No access -----
Dialup:  ----- No access -----           ----- No access -----
Remote:  ----- No access -----           ----- No access -----
Expiration:          (none)   Pwdminimum: 6   Login Fails: 771 1
Pwdlifetime:        180 00:00   Pwdchange: 14-JUN-1988 13:06
Last Login:          (none) (interactive), 9-JAN-1989 13:47 (non-interactive)
Maxjobs:             0   Fillm: 20   Byt1m: 4096
Maxacctjobs:         0   Shrfillm: 0   Pbyt1m: 0
Maxdetach:           0   BI01m: 18   JTquota: 1024
Prclm:               2   DI01m: 18   WSdef: 300
Prio:                4   AST1m: 24   WSquo: 512
Queprio:             0   TQElm: 10   WSextent: 1024
CPU:                 (none) Enqlm: 30   Pgflquo: 10000
Authorized Privileges:
  TMPMBX NETMBX
Default Privileges:
```

(continued on next page)

Example 3–1 (Cont.) Authorize Utility Example

```
TMPMBX NETMBX  
UAF> EXIT
```

Breakout Boxes

Overview

Breakout boxes are primarily used to isolate (or "break out") the signaling between two devices. By isolating the signals in this way, you can determine what signals are present, from which device, and in what order. This information can be very important when troubleshooting point-to-point problems.

Breakout boxes vary widely, but generally, they all perform the same basic functions.

Required Resources & Privileges

The breakout box is a hardware device that requires no privileges or special resources to use.

How to Use

Breakout boxes are best used when placed between a DCE and a DTE. It may be necessary to move the breakout box from one end of a cable to the other end of a cable to determine if the cable is the source of the problem.

Recommended Use

You can use a breakout box to examine the signals between a DTE and a DCE, or between two DCEs. Breakout boxes are especially useful in troubleshooting dialup problems, where signaling is critical. You can also use breakout boxes to diagnose cable problems and modem problems such as modem signaling problems.

DEC Extended LAN Management Software

Overview

DEC Extended LAN Management Software (DECelms) is a VMS layered product that allows users logged in to a VMS host to control and monitor any LAN bridge and FDDI wiring concentrator in an extended local area network. In addition, DECelms provides statistics gathered by the bridge to help monitor and troubleshoot the extended LAN.

A LAN bridge can serve either of two functions: it can be used as a LAN Traffic Monitor, or it can be the primary building block of an extended LAN. As a building block of an extended LAN, a LAN bridge is an intelligent device that acts as a store-and-forward station between two LAN segments. It operates at the data link level and is transparent to all other LAN stations.

The DEC Extended LAN Management Software resides on a VAX host. Corresponding management firmware resides in the LAN bridge. DECelms uses bridge management protocol for communications between the VAX host and the target LAN bridge.

DECelms allows you to do the following:

- Control and monitor the bridge management directory, bridges, bridge links, and the bridge-forwarding databases

Network Management and Troubleshooting Tools

DEC Extended LAN Management Software

- Display bridge and data-link counters, status, and characteristics
- Display bridge-forwarding databases
- Modify bridge parameters such as operational state, forwarding database entries, and spanning tree characteristics
- Invoke the bridge self-test
- Associate ASCII names with specific bridge physical addresses to make it easier to use DECelms management commands
- Perform protocol filtering for LAN Bridge 200s
- Perform LAN monitoring (for example, utilization percent) for LAN Bridge 200s

Required Resources & Privileges

To run DECelms, you need TMPMBX privilege on the VMS system. In addition, your system may also have an access control list (ACL) that specifies the users who can use DECelms and lists their privileges. For more information about ACLs, see the *DECelms Installation* manual and the VMS documentation.

Use of DECelms can be further restricted with DECelms privilege modes and bridge passwords. DECelms modes include the following:

- Nonprivileged mode, which limits DECelms use to monitoring device status and controlling the operation of DECelms itself. Nonprivileged users can use the SHOW, LIST, and MONITOR commands, and commands that control the DECelms registry and DECelms itself.
- Privileged mode, which allows users to control device configuration and operation. Privileged users have full use of all DECelms commands.

Your DECelms privilege mode is determined by the privilege associated with your DECelms process. To be a privileged user, you must have the ELMS\$RIGHTS_ID rights identifier associated with your DECelms process, in addition to the TMPMBX privilege. See the *DECelms Installation* manual for further information on how to add the ELMS\$RIGHTS_ID rights identifier to the system, and grant it to users.

How to Use

To run DECelms, enter the following command:

```
$ ELMS_LAN_MANAGEMENT
```

To get online help while using DECelms, enter the following command:

```
ELMS> HELP
```

To exit DECelms, enter the following command:

```
ELMS> EXIT
```

Recommended Use

Use DECelms to do the following:

- Monitor bridge operations
- Monitor traffic load by LAN segment
- Monitor traffic flow between segments
- Monitor error rates per segment

Network Management and Troubleshooting Tools DEC Extended LAN Management Software

- For network problems, examine bridge counters to determine which LAN segment may be causing the problem
- Isolate certain traffic (using destination addresses) to specific LAN segments
- Isolate multicasts to a specific LAN
- Locate nodes
- Check password protection on LAN Bridge 150s and LAN Bridge 200s

For example, for LAN segment communication problems, you can use DECelms to isolate the problem to the segment with the problem. To do this, use DECelms to poll the bridge and display the segment status. The display shows that a segmentation problem exists.

DECmcc Management Station for ULTRIX

Overview

DECmcc Management Station for ULTRIX is an ULTRIX-based network management tool for both IP and DECnet networks. DECmcc Management Station for ULTRIX collects, formats, and monitors network data from SNMP-based agents and DECnet nodes. You can display the data graphically and generate hardcopy reports. DECmcc Management Station for ULTRIX displays information using a DECwindows iconic map, which signals changes on network by changing the colors of network entities displayed. Using DECmcc Management Station for ULTRIX, you can observe trends, and take action to avoid problems.

DECmcc Management Station for ULTRIX allows you to do the following:

- Plot statistics for any IP host on the network
- Store historical data in an SQL database for later use in generating ASCII and graphical reports
- Store administrative reference information in an SQL database
- Display a map of IP routers on the network
- Poll IP hosts and DECnet nodes periodically to check their reachability
- Perform real-time monitoring of Ethernet traffic for the most active protocols and most active hosts
- Use a step-by-step interactive procedure to add private vendor Management Information Base (MIB) extensions
- Distribute map and polling operations on different ULTRIX hosts
- Send the results of polling and SNMP traps to an alarm window for display
- Trace the route of a packet from an IP source host to a destination host

DECmcc Management Station for ULTRIX consists of the following software:

- DECmcc Management Station for ULTRIX software running on a host running ULTRIX V4.0 or higher software
- An ULTRIX/SQL database running in the background as a data repository
- SNMP agent software for the entities to be managed which support SNMP

Network Management and Troubleshooting Tools

DECmcc Management Station for ULTRIX

Required Resources & Privileges

To run DECmcc Management Station for ULTRIX requires ULTRIX V4.0 or higher, ULTRIX/SQL (which comes bundled with ULTRIX V4.0), and DECnet to monitor DECnet nodes.

You need access to the superuser account to run DECmcc Management Station for ULTRIX.

How to Use

To run DECmcc Management Station for ULTRIX, the ULTRIX/SQL database must be running. In addition, you must perform various setup tasks, as described in the *DECmcc Management Station for ULTRIX Use* manual.

Recommended Use

DECmcc Management Station for ULTRIX is helpful in testing the "Connection Timed Out," "Host is Unreachable," and "Network is unreachable" problems. You can also use DECmcc Management Station for ULTRIX in troubleshooting "Broadcast Storm" and "LAN Segment Communication Problems."

LAN Traffic Monitor

Overview

The LAN Traffic Monitor (LTM) is a VMS layered product consisting of software and hardware that captures and presents traffic data for an Ethernet. The traffic data is available to multiple users on different nodes on the Ethernet.

The LAN Traffic Monitor allows you to actively monitor Ethernet usage by providing real-time data on Ethernet throughput and utilization. The LAN Traffic Monitor can collect data on all protocol types including, LAT, DECnet, TCP/IP, and Systems Communication Architecture (SCA), as well as 802.3 packets.

The hardware portion of the LAN Traffic Monitor is a LAN Bridge 100 (Rev. E, or higher) or a LAN Bridge 150 attached to the Ethernet cable. The LAN bridge runs monitoring software, and transmits information to a VMS layered application software program located on any VAX computer in the extended LAN.

The software portion of the LAN Traffic Monitor consists of two parts:

- LTM Listener software
- LTM User Interface

Using a LAN bridge as a monitoring device, the LTM Listener software counts and classifies Ethernet traffic, and periodically reports traffic statistics to the LTM User Interface. The LTM Listener can collect traffic data for time intervals from 1 to 48 hours. You can display LAN utilization data based on the traffic data collected for these periods.

The LTM User Interface collects and displays data received from the LTM Listener. By default, the LTM User Interface receives unsolicited Ethernet datagrams generated by the LTM Listener. However, you can request additional data not found in the unsolicited messages.

The LTM User Interface performs data reduction and presents statistics such as the following:

- Network traffic
- Top 10 talkers (current, long term)
- Throughput and utilization (current, long term, peak)

Network Management and Troubleshooting Tools LAN Traffic Monitor

- List of nodes on the extended Ethernet
- List of multicast addresses on the extended Ethernet
- List of nodes using a particular protocol
- Node traffic by protocol type
- Multicast traffic by protocol type
- Protocol type traffic

To monitor multiple Ethernet segments, you can distribute multiple LTM Listeners throughout the network. You can also distribute multiple User Interfaces throughout the extended LAN.

Required Resources & Privileges

To run LTM requires DECnet, a LAN Bridge 100 bridge (Rev. E or higher) or a LAN Bridge 150, and a REGIS terminal.

To use LTM, you need the following privileges:

- OPER
- NETMBX
- TMPMBX

How to Use

Before you run LTM, see the step-by-step procedure for troubleshooting "LAN Bridge Cannot Downline Load" in Chapter 5. This procedure explains how to evaluate the LAN bridge indicator lights to see if the bridge is set up to operate as a LAN Traffic Monitor.

After the bridge is properly set up, downline load the LTM software to the bridge.

To run LTM, enter the following command:

```
$ TRAFFIC_MONITOR/LAN
```

Or you can define LTM as the symbol for TRAFFIC_MONITOR/LAN, as follows:

```
$ LTM = TRAFFIC_MONITOR/LAN
```

Now, to run LTM use the symbol name as follows:

```
$ LTM
```

To get online help while using LTM, enter the following command:

```
LTM> HELP
```

To exit LTM, press the PF4 key:

```
LTM> PF4
```

Recommended Use

Use the utilization information that LTM generates to help determine LAN traffic patterns and identify possible performance problems. You can also use LTM to help isolate network problems to the node or protocol causing the problem.

For the babbling device problem, use LTM to display traffic rates. If the babbling device is not transmitting corrupt information for its own address, the LTM display lists the babbling device as the top talker on the segment.

Network Management and Troubleshooting Tools

LAN Traffic Monitor

Example 3–2 LAN Traffic Monitor Example

LAN Traffic Monitor V1.1.0 26-JAN-1989 16:14:58 Listener uptime 03 06:55:31

Current Top Ten Talkers Display

Total Node Addresses : 125 Report Interval : 2.95

Address	Name	Count	Frames/Sec	% Total
AA-00-04-00-3F-12	NODEA	393	133.2	77.7 % ❶
08-00-2B-06-51-C9	LAT_1	25	8.5	4.9 %
AA-00-04-00-4E-13	NODED	23	7.8	4.5 %
08-00-2B-05-E2-63	LAT_2	19	6.4	3.7 %
AA-00-04-00-40-12	NODEM	18	6.1	3.6 %
AA-00-04-00-EE-11	NODEP	15	5.1	3.0 %
AA-00-04-00-7A-11	NODEB	10	3.4	2.0 %
<Other nodes>		3	1.0	0.6 %
<Total>		506	171.5	100.0 %

Generally, any node generating greater than 50% of the traffic could be a babbling device. In Example 3–2, NODEA is probably the babbling device as it is generating more than three-quarters of the network traffic. ❶

LAT Control Program

Overview

The LAT Control Program, which is bundled with the VMS and ULTRIX operating systems, allows you to configure and control the Local Area Transport (LAT) protocol on VMS and ULTRIX host systems. The LAT Control Program is known as LATCP on VMS systems, and lcp on ULTRIX systems. VMS and ULTRIX systems that join a LAT configuration are called **service nodes**. You can use the LAT Control Program to do the following:

- Start and stop the LAT port driver (LTDRIVER) on VMS systems
- Start and stop LAT service on ULTRIX systems
- Specify operational characteristics for your service node and its services
- Display the status for your LAT service node
- Show and zero LAT counters

Required Resources & Privileges

On VMS Systems

Running LATCP requires the following privileges:

- CMKRNL
- TMPMBX
- NETMBX

On ULTRIX Systems

To run lcp on an ULTRIX system, you must have read and write access to a terminal.

How to Use

On VMS Systems

To run LATCP, enter the following command:

```
$ MCR LATCP
```

To get online help while using LATCP, enter the following command:

```
LCP> HELP
```

To exit LATCP, enter the following command:

```
LCP> EXIT
```

On ULTRIX Systems

Use the following command syntax to run the LAT Control Program on an ULTRIX system:

```
/etc/lcp [options]
```

The LAT control program executes the command you specify, and returns to the default prompt in multiuser mode.

To get help on the LAT control program, see the *ULTRIX-32 Reference Pages, Section 8* or the online reference page for lcp.

Recommended Use

Use LATCP to display and define port characteristics; display LAT counters, group code definitions, and servers connected; and to stop the LAT protocol. You can use LATCP to help isolate the source of the problem "Service Unavailable" by verifying that appropriate group codes are established and that LAT is running (Status: Active).

Example 3-3 uses LATCP to gather information on the VMS host node BOSTON. (Use this in conjunction with TSM.)

Example 3-3 VMS LATCP Example

```
$ MCR LATCP
LCP> SHOW CHARACTERISTICS
LCP Characteristics

Node name = \BOSTON\
Node Identification = \Unauthorized access is prohibited.\
Groups = (0,40)
Multicast timer = 60 seconds
LAT Version = 5.1           LAT Protocol is active

Service Names and Ids:

Service name : \BOSTON\           rating : <auto>
              ID : \Unauthorized access is prohibited.\

Node Links:

Link name = \LAT$LINK\
Link device = \XQA0:\
Groups = ()
Link-specific services:
Status = Active
```

(continued on next page)

Network Management and Troubleshooting Tools

LAT Control Program

Example 3–3 (Cont.) VMS LATCP Example

```
LCP> EXIT
```

Example 3–4 shows how to display LAT characteristics on an ULTRIX host.

Example 3–4 ULTRIX lcp Example

```
% /etc/lcp -d
Node name/identification: ALARMS / ULTRIX-32
Service name/identification: ALARMS / ULTRIX LAT SERVICE
Groups: 40 44
Multicast timer: 30 seconds
LAT version: 5 eco: 1 LAT Protocol is active
%
```

netstat command

Overview

The ULTRIX command, `netstat`, gives the local host's status with respect to the network by symbolically displaying the contents of network-related data structures. You can select several forms of display, each allowing you to specify a particular type of information to emphasize, including the following:

- A list of active sockets for each protocol
- The contents of one of the other network data structures, according to the option selected
- The information regarding packet traffic on the configured network interfaces, displayed at the interval you specify
- The state of all active sockets from those using any of the protocols listed in the `/etc/protocols` file
- The information for the memory management routines
- The statistics per protocol

Required Resources & Privileges

None.

How to Use

Use the following command syntax to display network-related information:

```
# netstat [options...]
```

Table 3–2 shows the netstat command options.

Table 3–2 ULTRIX netstat Command Options

Option	Information Displayed
-A	Address of any associated protocol control blocks
-a	Information for all sockets
-f { <i>address_family</i> }	Statistics or address control block reports for the specified address family
-h	State of the IMP host table
-I { <i>interface</i> }	Information about the specified interface only
-i	Status information for autoconfigured interfaces
-m	Information for the memory management routines
-n	Network addresses in number form rather than symbolic form
-r	Routing tables
-s	Statistics per protocol
-t	Time until interface watchdog routine starts

Recommended Use

The netstat -i command gives statistics on each active or open network interface. Outgoing packet errors (Oerrs) indicate that the local host may have a problem. Incoming errors (Ierrs) indicate a potential problem with the network connected to the interface. Example 3–5 shows a normal display (no Ierrs or Oerrs) from the netstat -i command.

Example 3–5 netstat -i Example

```
% netstat -i
Name Mtu Network Address Ipkts Ierrs Opkts Oerrs Coll
qe0 1500 DECnet RTR 8324125 0 8347463 2 237706
qe0 1500 16.31.16 RTR.xxx.corp.co 8324125 0 8347463 2 237706
dmv0 1284 16.10.16 xx2nnn.pa.corp. 10746232 0 9569078 95 62
dmv1 1284 16.10.16 xx2yyy.pa.corp.c 4880317 0 4570916 0 0
dmv2 1284 16.10.16 RTR2zso 1368208 0 1729512 0 31
lo0 1536 loop localhost 909234 0 909234 0 0
idl0 549 ptopnet 128.45.10.131 0 0 15 0 0
idl1* 549 none none 0 0 0 0 0
idl2* 549 none none 0 0 0 0 0
idl3* 549 none none 0 0 0 0 0
idl4 549 128.45.110 128.45.110.111 0 0 15 0 0
%
```

Network Control Program

Overview

The Network Control Program (NCP), bundled with DECnet software, runs on both the VMS and ULTRIX operating systems, and allows you to configure and control DECnet networks. Using NCP, you can monitor network resources, test network components by manipulating the configuration database, and display error counter information.

Network Management and Troubleshooting Tools

Network Control Program

Each node (or host) in a DECnet network has a configuration database that consists of the following databases:

- Node database with a record for each node in the network, including the local node
- Circuit database with a record for each circuit known to the local node
- Line database with a record for each physical line known to the local node
- Logging database with a record for each sink (logged events are sent to the sinks)

For DECnet-VAX networks, the Network Ancillary Control Process (NETACP) provides a default object database with a record for each object known to the network, including objects that are defined when you start the local node, such as the File Access Listener (FAL).

The DECnet configuration database consists of two distinct databases, one volatile and one permanent. If VAX PSI is included in the network, a VAX PSI configuration database consisting of a volatile database and a permanent database exists at the local DTE.

NCP commands allow you to create, modify, clear, and display parameters for the configuration database, including the following:

- Nodes
- Circuits
- Lines
- Objects
- Logging

For DECnet-VAX networks, NCP commands allow you to create, modify, clear, and display parameters for the configuration database, including the following:

- Routing
- Logical links
- X.25 protocol modules

Network Management and Troubleshooting Tools Network Control Program

- X.25 or X.29 server modules
- X.25 access modules
- Network access control

NCP commands also allow you to display and clear counters for circuits, lines, the executor, and other nodes.

Required Resources & Privileges

DECnet-VAX Requirements

To run NCP on a DECnet-VAX network, you need the following privileges to display information from the NCP databases:

- NETMBX
- TMPMBX

To change NCP information, you need the following privileges:

- OPER privilege to modify the volatile database
- SYSPRV privilege to modify the permanent database

DECnet-ULTRIX Requirements

To run NCP on a DECnet-ULTRIX network requires the following:

- No special privileges are required to use the NCP show or list commands to display information from the databases,
- Log in to the super user account to use an NCP command that modifies a database.

How to Use

DECnet-VAX

To run NCP on a DECnet-VAX network, enter the following command:

```
$ MCR NCP
```

NCP uses different commands for manipulating the volatile and permanent databases. When you change the NCP databases, be sure to use the proper commands, as shown below:

- For the volatile databases, use the SET and SHOW commands.
- For the permanent databases, use the DEFINE and LIST commands.

To perform NCP operations on remote nodes, do one of the following:

- Use the NCP command, SET EXECUTOR, to change your executor from the current (local) executor to the remote executor, then issue NCP commands on that executor.
- Use the following command prefix to direct a remote node to execute an NCP command:

```
NCP> TELL remote-node-id [NCP command]
```

To get online help while using NCP, enter the following command:

```
NCP> HELP
```

To exit NCP, enter the following command:

```
NCP> EXIT
```

Network Management and Troubleshooting Tools

Network Control Program

DECnet-ULTRIX

Use one of the following commands to run NCP:

```
% ncp
```

```
% ncp command
```

(*command* is any ncp command.)

```
% ncp < filename
```

(*filename* is a shell script that contains a sequence of ncp commands.)

To exit ncp, enter one of the following commands:

```
ncp>exit
```

```
ncp>quit
```

```
ncp>ctrl/d
```

To get information on ncp commands while using ncp, enter following command:

```
ncp>help
```

Recommended Use

Use NCP to do the following:

- Display, modify, and delete information about the local or remote nodes, lines, and circuits
- Load unattended systems
- Test network components.

Most of the step-by-step procedures in Chapter 5 use NCP to help gather information about a problem or to correct a problem.

To troubleshoot most "Network Object Unknown" problems, use NCP to display object characteristics, as shown in Example 3-6.

Example 3-6 VMS Network Control Program Example

```
$ MCR NCP
NCP> SHOW KNOWN OBJECT CHARACTERISTICS
Known Object Volatile Characteristics as of  9-JAN-1989 13:52:05

Object = TASK
Number           = 0
User id          = BADACCOUNT
Password         = PURPOSEFULLY

Object = FAL
Number           = 17
File id          = FAL.EXE

Object = HLD
Number           = 18

Object = NML
```

(continued on next page)

Network Management and Troubleshooting Tools Network Control Program

Example 3–6 (Cont.) VMS Network Control Program Example

```
Number          = 19
File id         = NML.EXE

Object = REMACP

Number          = 23
Process id      = 0000002B

Object = MIRROR

Number          = 25

Object = EVL

Number          = 26
Process id      = 00000250

Object = MAIL

Number          = 27
File id         = MAIL_SERVER.EXE
Proxy access    = outgoing
Alias outgoing  = Enabled

Object = PHONE

Number          = 29
File id         = PHONE.EXE
Proxy access    = outgoing
Alias incoming  = Disabled

Object = CTERM

Number          = 42
Process id      = 0000002B

NCP> EXIT
```

Network Management and Troubleshooting Tools

Network Control Program

Example 3-7 gives an example of the display for an ULTRIX operating system.

Example 3-7 ULTRIX Network Control Program Example

```
ncp> show known objects characteristics
Known Object Volatile Characteristics as of Wed Apr 18 16:03:26 EDT 1990

Object = TELL
Number           = 0
File             = /usr/bin/tell
Type             = Sequenced Packet
Accept           = Immediate

Object = DEFAULT
Number           = 0
Type             = Sequenced Packet
Accept           = Immediate

Object = fal
Number           = 17
File             = /usr/etc/fal
Default User     = guest
Type             = Sequenced Packet
Accept           = Deferred

Object = nml
Number           = 19
File             = /usr/etc/nml
Default User     = guest
Type             = Sequenced Packet
Accept           = Deferred

Object = mir
Number           = 25
File             = /usr/etc/mir
Default User     = guest
Type             = Sequenced Packet
Accept           = Deferred

Object = mail11
Number           = 27
File             = /usr/etc/mail11dv3
Default User     = daemon
Type             = Sequenced Packet
Accept           = Deferred

ncp>
```

NMCC/DECnet Monitor

Overview

NMCC/DECnet Monitor is a layered application that runs on the VMS operating system. NMCC/DECnet Monitor collects, formats, and highlights network data for Phase III and Phase IV DECnet network nodes. You can graphically display the resulting data or generate a printed report.

Using NMCC/DECnet Monitor, you can observe trends in the network, and take action to avoid problems before they disrupt network service.

Network Management and Troubleshooting Tools NMCC/DECnet Monitor

NMCC/DECnet Monitor consists of the following software modules:

- Kernel
- User Interface subsystem
- NMCC/DECnet Reports subsystem

The Kernel is a data acquisition process that collects the following information through polling and event logging:

- Status information for nodes, lines, and circuits
- Characteristics information for nodes, lines, and circuits
- Traffic counters
- Error counters

You determine the nodes that the Kernel polls, the polling rates, and the nodes that log events.

The User Interface subsystem provides the command interface to the NMCC/DECnet Monitor. Multiple users can monitor different network components simultaneously and gather information about the network. NMCC/DECnet Monitor graphically displays the following types of information:

- Selected error statistics
- Network traffic statistics
- Network status information

The User Interface subsystem also allows you to build and manipulate databases maintained by NMCC/DECnet Monitor, including a network component reference database. This database contains a description of all nodes and physical links in the network. It contains information about the CPU type and the operating system type and version.

The NMCC/DECnet Reports subsystem uses the log file of statistical data to produce network reports. The NMCC/DECnet Reports subsystem provides error, traffic, and configuration reports.

Network Management and Troubleshooting Tools

NMCC/DECnet Monitor

Required Resources & Privileges

To run NMCC/DECnet Monitor requires DECnet, REGIS, and the run-time options for the GKS and RdB software. DECnet routing on the kernel node is recommended, but not required.

You need the following privileges to run NMCC/DECnet Monitor:

- DETACH
- NETMBX
- SYSNAM
- TMPMBX

How to Use

To run NMCC/DECnet Monitor, enter the following command:

```
$ NMCC [/KERNEL=kernel_system]
```

To get help while using NMCC/DECnet Monitor, do one of the following:

- Enter the HELP command.
- Press the Help key during any operation.
- Press the Help key while the cursor is on a particular field in the display to get information on that field.

To exit NMCC/DECnet Monitor, type EXIT or press Ctrl/Z.

Recommended Use

NMCC/DECnet Monitor is especially helpful in troubleshooting the "Remote Node Is Not Currently Reachable" problem, and the "Partitioned Area" problem. If the problem occurs within the defined NMCC/DECnet Monitor database, both problems show up on the NMCC/DECnet Monitor display as circuit or router problems. See the *NMCC/DECnet Monitor User's Guide* for a detailed example. You can also use NMCC/DECnet Monitor to periodically poll the network. This allows you to identify potential problems before they affect users.

NMCC/VAX ETHERnim

Overview

NMCC/VAX ETHERnim (network integrity monitor) is a layered application that runs on VMS systems and allows you to manage Ethernet local area networks. NMCC/VAX ETHERnim gathers information about Ethernet nodes, verifies that nodes are reachable, graphically displays the local area network topology, and monitors Ethernet traffic.

Using NMCC/VAX ETHERnim, you can observe trends in the network, and take action to avoid problems before they disrupt network service.

Specifically, NMCC/VAX ETHERnim allows you to do the following:

- Test the Ethernet nodes to help isolate and identify problems
- Build a permanent database containing information about each node directly connected to the Ethernet local area network, including the following:
 - Ethernet, DECnet, and hardware addresses
 - Ethernet and other communications port types
 - DECnet ID messages, management versions, router versions and types, and NSP versions

Network Management and Troubleshooting Tools NMCC/VAX ETHERnim

- Administrative reference data
- Monitor the following kinds of network activity:
 - Nodes transmitting the most messages
 - Nodes receiving the most messages
 - Traffic to and from a specified node
 - Nodes transmitting multicast messages
 - Protocol types currently in use
 - Nodes that are loading the network and slowing down the response of other nodes
- Examine user's nodes for changes to the hardware or software configuration
- Manually construct and maintain a topological representation of the network, including transparent network devices, such as the DEMPR, DELNI, and DEREPI

If you install remote test command files on VMS target nodes, NMCC/VAX ETHERnim provides additional capabilities. For example, you can do the following for VMS target systems:

- Test the communications path to each node up to the DNA application layer by performing a loopback test to a specified user account
- Perform information access to DECnet nodes on the Ethernet to determine the remote node's system type and hardware and software configuration

Required Resources & Privileges

To run NMCC/VAX ETHERnim requires DECnet and REGIS software.

To run NMCC/VAX ETHERnim, you need the following privileges:

- DIAGNOSE
- LOG_IO
- NETMBX
- PHY_IO
- TMPMBX

How to Use

To run NMCC/VAX ETHERnim, enter the following command:

```
$ ETHERNIM
```

NMCC/VAX ETHERnim provides the following types of online help:

- Introductory help gives a brief overview of NMCC/VAX ETHERnim before startup.
- Command help (available through the main menu) gives either brief or complete information on a topic.

To display a one-line HELP message on the status line, press the PF2 key once. To display a full screen of additional information, press the PF2 key twice.

- Tutorial help (available through the main menu) provides information on network management concepts using NMCC/VAX ETHERnim.

Network Management and Troubleshooting Tools

NMCC/VAX ETHERnim

To display the tutorial information, use the main menu HELP command.

- Template help provides forms that enable you to easily use the SHOW, VERIFY PATH, TEST NODE, and EDIT commands.

To exit NMCC/VAX ETHERnim, enter the following command:

Command: EXIT

Recommended Use

Use NMCC/VAX ETHERnim to test the path at various DECnet layers, and to test for node reachability. This is especially helpful with the "Line Synchronization Lost" problem on Ethernet circuits. See the *NMCC/VAX ETHERnim User's Guide* for a detailed example.

You can also use NMCC/VAX ETHERnim to periodically poll the network. This allows you to identify potential problems before they affect users.

ping command

Overview

The ping command in ULTRIX allows you to determine the reachability of network hosts on an internet.

The ping command performs an ICMP echo request to the hostname specified. When the request is successful, the remote host sends the data back to the local host.

Required Resources & Privileges

None.

How to Use

Use the following command syntax with the ping command:

```
# /etc/ping [options] hostname
```

Table 3-3 shows the command options.

Table 3-3 Options for the ping Command

Option	Function
-l	Displays a long version of the ping results.
-v	Displays a verbose version of the ping results, similar to the long version, and including ICMP packets other than ECHO RESPONSE packets.
-r	Executes the command for a host directly connected to the local host. With this option, the ping command bypasses normal routing tables and sends the request directly to a host on an attached network. If the host is not on a directly-attached network, the local host receives an error message.

To terminate the ping command output, press Ctrl/C. When terminated, the ping command displays statistics on packets sent, packets received, the percentage of packets lost, and the minimum, average, and maximum round trip packet times.

Recommended Use

Use the ping command for direct and indirect routing problems such as "Host is unreachable," "Connection timed out," and "Network is unreachable."

When using the ping command for fault isolation, first test the local host to verify that the local host is running. If the local host returns the data correctly, use the ping command to test remote hosts further and further away from the local host.

If you do not specify command options, the remote host returns the message, "hostname is alive." If the remote host does not respond to the request, the ping command displays the message, "no answer from hostname."

In the verbose or long form, the ping command displays each ICMP request in sequence, the number of bytes received from the remote host, and the round trip time on a per request basis. Example 3–8 shows the results of a ping command to host hostb.

Example 3–8 Long Output from the ping Command

```
PING hostb.corp.com (16.20.32.2): 56 data bytes
64 bytes from 16.20.32.2: icmp_seq=0. time=20. ms
64 bytes from 16.20.32.2: icmp_seq=1. time=10. ms
64 bytes from 16.20.32.2: icmp_seq=2. time=10. ms
64 bytes from 16.20.32.2: icmp_seq=3. time=20. ms
64 bytes from 16.20.32.2: icmp_seq=4. time=10. ms
64 bytes from 16.20.32.2: icmp_seq=5. time=10. ms
64 bytes from 16.20.32.2: icmp_seq=6. time=10. ms
64 bytes from 16.20.32.2: icmp_seq=7. time=20. ms
64 bytes from 16.20.32.2: icmp_seq=8. time=10. ms
```

Protocol Analyzers

Overview

Most protocol analyzers allow you to view the frame format of data at the physical and data link layers for many different protocols, and determine if data has been sent correctly. Some protocol analyzers allow you to view data at higher protocol levels. You can view the frame formats on a frame-by-frame basis, and use this information to determine whether the the problem is a transmission or reception problem.

Required Resources & Privileges

Protocol analyzers are hardware devices that requires no privileges or special resources to use.

How to Use

See the instructions for your specific protocol analyzer.

Recommended Use

Protocol analyzers are useful in solving network problems where it is difficult to determine the cause of the problem, as in the following examples:

- File transfer problems resulting in the receipt of corrupted data
A protocol analyzer can help you determine which end of the communications link is corrupting the data.

Network Management and Troubleshooting Tools

Protocol Analyzers

- RS232 handshake problems

This case involves modem communication problems, where the modem lights are on, but the line is not synchronizing and the RS232 signals are dropping. In this case, you use the protocol analyzer to correlate the signals with the data loss.

- ICMP redirect messages

By evaluating the frame format, you can determine if the IP router is sending the correct redirect information by comparing the redirect information with routing information in routing tables.

syslog daemon

Overview

The ULTRIX syslog daemon records system messages into a set of files.

The syslog daemon starts from the `/etc/rc.local` file when you boot the system, and whenever it receives a hangup signal. Before the syslog daemon starts logging system messages, it scans the `/etc/syslog.conf` file to determine its configuration information. The configuration information determines the files into which the syslog daemon logs system messages.

System messages can contain a priority code indicating the type and severity of the message. For example, system messages can indicate error conditions and warnings.

The syslog daemon is available only to user-level programs (daemons) and cannot record kernel-based error messages. Kernel-based errors are logged and translated through the error logger and uerf.

For a complete description of the syslog daemon see the *ULTRIX-32 Reference Pages*, Section 8.

Required Resources & Privileges

To use the information collected by the syslog daemon requires no special privileges.

How to Use

To review the log files generated by the syslog daemon, do the following:

1. Change your current directory to the `/etc` directory using the following command:

```
# cd /etc
```

2. Use the `cat` command to display the contents of the `syslog.conf` file, which tells you where the syslog files are kept on your system:

```
# cat syslog.conf
```

3. Change your current directory to the directory specified in the `syslog.conf` file. In the following example, the syslog directory is `/usr/spool/mqueue`.

```
# cd /usr/spool/mqueue
```

4. Display the list of available syslog files:

```
# ls
```

Network Management and Troubleshooting Tools syslog daemon

5. Use the `cat` command to display the contents of the syslog file you want to see. In the following example, the file is `syslog.5`.

```
# cat syslog.5
```

Recommended Use

Use the syslog daemon to help solve session layer problems such as access control problems for DECnet and IP.

Example 3–9 shows typical syslog output, including reference to a potential DECnet problem ❶, "Insufficient Network Resources."

Example 3–9 syslog Example

```
Jun 22 04:05:02 syslog restart
Jun 22 04:08:15 localhost: 20287 sendmail: AA20287:
message-id=<9006220808.AA20287@bambam.corp.com>
Jun 22 04:08:15 localhost: 20287 sendmail: AA20287: from=root, size=301,
class=0
Jun 22 04:08:17 localhost: 20290 sendmail: AA20287: to=postmaster,
delay=00:00:03, stat=Sent
Jun 22 04:12:25 localhost: 20307 sendmail: AA12945: to=hocus::bill,
delay=1+14:26:30, stat=Deferred: ❶Insufficient Network Resources at node
hocus
Jun 22 04:33:45 localhost: 20435 sendmail: AA20435:
message-id=<9006220833.AA20435@bambam.corp.com>
Jun 22 04:33:45 localhost: 20435 sendmail: AA20435:
from=<bedrck::alice>,
size=531, class=0
Jun 22 04:33:47 localhost: 20438 sendmail: AA20435: to=<ida>,
delay=00:00:12, stat=Sent
Jun 22 04:36:31 localhost: 20451 named-xfer: bad response to SOA query from
98.0.4.13, zone auth.corp.com: rcode 14, aa 0, ancourt 1006632960, auctount
808599092
Jun 22 04:36:31 localhost: 213 named: zoneref: Masters for secondary zone
auth.corp.com unreachable
Jun 22 05:08:18 localhost: 219 mountd: couldn't reply to rpc call
Jun 22 05:08:23 localhost: 219 mountd: couldn't reply to rpc call
Jun 22 05:09:31 localhost: 20733 dlogind: connect from ESSEX::JLS
Jun 22 05:10:06 localhost: 20733 dlogind: disconnect from ESSEX::JLS
Jun 22 05:12:25 localhost: 20805 sendmail: AA12945: to=hocus::bill,
delay=1+15:26:30, stat=Deferred: Insufficient Network Resources at node hocus
Jun 22 05:13:28 syslog: shutdown within 30 seconds
Jun 22 05:13:28 localhost: 156 dnet_spawner: Received signal number 15, exiting.
Jun 22 05:13:28 localhost: 344 ntpd: terminated: (sig 15)
Jun 22 05:13:30 localhost: 417 snmpd: Exit snmpd at Fri Jun 22 05:13:28 1990 ...
Jun 22 10:34:51 localhost: 502 ntpd: /usr/etc/ntpd version $Revision: 3.4.1.9 $
Jun 22 10:34:51 localhost: 502 ntpd: patchlevel 13
Jun 22 10:34:52 localhost: 504 sendmail: alias database rebuilt by root
Jun 22 10:34:53 localhost: 505 timed: THIS MACHINE IS A SLAVE
Jun 22 10:34:53 localhost: 504 sendmail: 11 aliases, longest 30 bytes, 289
bytes total
Jun 22 10:35:07 localhost: 550 snmpd: Start snmpd version 3.2 at Fri Jun 22
10:35:06 1990
```

Network Management and Troubleshooting Tools

Terminal Server Manager Software

Terminal Server Manager Software

Overview

Terminal Server Manager (TSM) software is a VMS layered product that allows you to control and observe terminal servers anywhere within an extended local area network.

TSM runs from a central location on a host system and allows you to do the following:

- Set up and manipulate a database of terminal servers
- Propagate server information from the TSM management directory to the DECnet database
- Set up and maintain terminal servers on the same local area network as the host system
- Perform troubleshooting on the terminal servers registered in the database of terminal servers
- Define groups of terminal servers, each of which you can manage as a single entity
- Test the reachability of defined terminal servers

Required Resources & Privileges

Running TSM requires DECnet software, read access to the TSM management directory file, and OPER privilege.

To set terminal server parameters with TSM also requires the following privileges and password:

- Privileges on the server
- Server access password
- OPER privilege
- NETMBX privilege
- TMPMBX privilege

How to Use

You can use the TSM software in two ways. First, you can run the TSM program and enter TSM commands at the TSM> prompt. Second, you can run the TSM program and enter TSM commands directly from the DCL prompt.

Before you run TSM, define TSM as the symbol for `TERMINAL_SERVER_MANAGER` as shown:

```
$ TSM = = TERMINAL_SERVER_MANAGER
```

Then, to run TSM, and enter commands from the DCL prompt, enter the following command:

```
$ TSM [TSM command]...
```

Network Management and Troubleshooting Tools

Terminal Server Manager Software

To get online help for TSM, enter the following command:

```
TSM> HELP
```

To exit TSM, enter the following command:

```
TSM> EXIT
```

Recommended Use

Use TSM to display group code definitions and port characteristics, and to define various parameters.

For example, with the problem, "Terminal Server Connection Failure", you can use TSM to display and modify the port characteristics. In Example 3–10, information displayed for port 4 shows that the required group code (group code 40) ❶ was not enabled on the user's port (port 4 in this case). Enabling port 4 solves the problem.

To verify the solution, have the user attempt to connect to node ABC again.

Example 3–10 Terminal Server Manager Example

```
$ TSM
TSM> USE SERVER abc
TSM> SET PRIVILEGE
password> xxxxxx
TSM> SHOW PORT 4
Port 3: r                               Server: abc
Character Size:          8               Input Speed:      9600
Flow Control:           XON              Output Speed:     9600
Parity:                 None             Modem Control:    Disabled
Access:                 Local            Local Switch:     None
Backwards Switch:      None              Name:             PORT_3
Break:                  Local            Session Limit:    4
Forwards Switch:       `                 Type:             Soft
Preferred Service: DELNI
Authorized Groups: 41- 44, 46- 49, 54, 56- 59, 79
(Current) Groups: 41- 44, 46- 49, 54, 56- 59, 79 ❶
Enabled Characteristics:
Autobaud, Autoprompt, Broadcast, Input Flow Control, Loss Notification,
Message Codes, Output Flow Control, Verification
TSM> DEFINE PORT 4 GROUP 40 ENABLED
TSM> SET PORT 4 GROUP 40 ENABLED
TSM> EXIT
```

traceroute command

Overview

The traceroute command in ULTRIX displays the route a packet takes to a network host.

Note

Traceroute does not currently ship with ULTRIX Version 4.0, and earlier versions of ULTRIX do not support it. However, traceroute is available on other UNIX-based systems. You can obtain a copy of the traceroute software through FTP from various sites. See Appendix B for information on using FTP.

Network Management and Troubleshooting Tools

traceroute command

The traceroute command sends User Datagram Protocol (UDP) packets (known as **probe packets**) to an unused port on the remote host, and listens for ICMP replies from IP routers. The probe packets are sent with a small "time to live" parameter (ttl), which specifies the maximum number of hops a packet can take to reach its destination.

The two ICMP messages of interest when using traceroute are "time exceeded" and "port unreachable." The ICMP message, "time exceeded" means that the IP router that received the probe packet cannot forward it any further due to the ttl value. The ICMP message, "port unreachable," means that the host that received the probe packet cannot access the port intended for the probe packet.

In displaying a routing path, traceroute starts by specifying a ttl of one, and increasing the ttl by one for each probe packet it sends. The "time exceeded" messages tells you which IP routers are processing the packets. The "port unreachable" message tells you that the probe packet reached its intended destination, but could not access the intended port.

Required Resources & Privileges

To use the traceroute command you must be logged in to the superuser account.

How to Use

Use the following syntax with the traceroute command:
`traceroute [options...] host [packetsize]`

Network Management and Troubleshooting Tools

traceroute command

Table 3–4 lists the traceroute command options.

Table 3–4 ULTRIX traceroute Command Options

Option	Meaning
-m <i>{max_ttl}</i>	Sets the maximum time-to-live (ttl) used in outgoing probe packets. The time-to-live parameter specifies the maximum number of hops a packet can take to reach its destination. The default is 30 hops.
-n	Displays hop addresses numerically only, rather than both numerically and symbolically.
-p <i>{port}</i>	Sets the base UDP port number to be used in outgoing probe packets. The default is 33434. The traceroute command uses the port information to select an unused port range if a port in the default range is already used.
-r	Bypasses the normal routing tables and sends the probe packet directly to a host on an attached network. If the host is not on a directly-attached network, the traceroute command returns an error.
-s <i>{IP_address_number}</i>	Uses the specified IP address number as the source address in outgoing probe packets. On hosts with more than one IP address, the -s option forces the traceroute command to use the specified source address rather than any others the host might have. If the IP address is not one of the receiving host's interface addresses, the traceroute command returns an error and does not send a probe packet.
-t <i>{type-of-service} {value}</i>	Sets the type-of-service in probe packets to the specified value. The default is zero. The value must be a decimal integer in the range 0-255. The -t option tells you if different types of service result in different paths. This option is available only in Berkeley UNIX (4.4BSD) environments. Not all types of service are legal or meaningful. Useful values for this option are 16 (low delay) and 8 (high delay). See RFC 791, <i>Internet Protocol</i> for more information on types of service.
-v	Displays verbose output, which includes received ICMP packets other than "time exceeded" and "port unreachable."
-w <i>{wait_time}</i>	Sets the time (in seconds) to wait for a response to a probe. The default is three seconds.
packetsize	Sets the packet size (in bytes) for the probe packet. The default size is 38 bytes.

Network Management and Troubleshooting Tools

traceroute command

To use the traceroute command to display a route, begin sending probe packets with a time-to-live (ttl) parameter of 1. For each probe packet you send, increase the ttl parameter by one until you receive the ICMP message, "port unreachable."

The traceroute command sends three probe datagrams for each ttl setting, and displays a line showing the following:

- ttl
- Address of the IP router
- Round trip time of each probe datagram

If multiple IP routers respond to the probe, the traceroute command displays the address of each IP router. If the traceroute command does not elicit a response in three seconds (the default wait time), the traceroute command displays an asterisk (*) for the probe.

Recommended Use

Use the traceroute command for manual network testing, measurement, and management.

uerf command

Overview

The uerf command in ULTRIX runs the ULTRIX error report formatter, and displays the contents of system events recorded in the error log file, /usr/adm/syserr/syserr.hostname. The error log file is a data file that is not readable without use of the uerf command.

The events recorded in the /usr/adm/syserr/syserr.hostname file include error messages relating to the system hardware and the software kernel, as well as information about system status, startup, and diagnostics.

For a complete description of the uerf command see the *ULTRIX-32 Reference Pages*, Section 8. See the *ULTRIX-32 Guide to the Error Logger* for more information about the error logger.

Required Resources & Privileges

None.

How to Use

Use the following syntax with the uerf command:
% /etc/uerf [options...]

Table 3–5 lists the uerf command options.

Table 3–5 ULTRIX uerf Command Options

Option	Meaning
-A {adapter}	Displays information on adapter and controller errors. If you do not specify a type, the -A option records information for all adapter types.

(continued on next page)

Network Management and Troubleshooting Tools

uerf command

Table 3–5 (Cont.) ULTRIX uerf Command Options

Option	Meaning
-c <i>{classes}</i>	Displays information about classes of events, including the following: <ul style="list-style-type: none"> • err (hardware and software detected errors) • maint (events occurring during system maintenance) • oper (messages regarding system status, autoconfiguration, device status or error, time stamps, system startup and shutdown)
-D <i>{disks}</i>	Displays information for the disk device that you specify. You can specify a type (for example, ra60) or a class (for example, ra). If you do not specify a type, the -D option records information for all disk types.
-f <i>{filename}</i>	Displays information from the specified file rather than from the default error log file. You must specify the full path name for the file, and you cannot use the -n option with the -f option.
-h	Displays a brief help message about the uerf command.
-H <i>{host}</i>	Displays errors for the host that you specify. Use the -H option when you log errors from multiple remote hosts to a single error log file on the local host.
-M <i>{mainframe_errors}</i>	Displays mainframe errors according to the type that you specify. Types include cpu (cpu-related errors such as machine checks), and mem (memory-related errors such as single-bit CRD (corrected read data)) and double-bit uncorrectable errors.
-n	Displays errors as they occur in real time before logging them in the error log file. You cannot use the -n option with the -f option.
-o <i>{output_format_type}</i>	Displays errors in brief, full, or terse format. The default format is brief.
-O <i>{operating_system_events}</i>	Displays information on operating system events such as exceptions and faults. If you do not specify any events, the -O option records all operating system events.
-R <i>{reverse_chronological_order}</i>	Displays error information starting with the most recent information.
-r <i>{records}</i>	Displays information for the record code that you specify.
-s <i>{sequence_numbers}</i>	Displays errors for selected sequence numbers.
-T <i>{tapes}</i>	Displays information for the tape types that you specify (for example, tk50).
-t <i>{time_range}</i>	Displays errors for the time period you specify.

(continued on next page)

Network Management and Troubleshooting Tools

uerf command

Table 3–5 (Cont.) ULTRIX uerf Command Options

Option	Meaning
-x	Excludes the specified option from the report. The -x option does not affect the -f, -h, -H, -n, -o, -R, -t options.
-Z	Displays the error entry in hex form.

Recommended Use

Use the uerf command to diagnose kernel and hardware errors on ULTRIX systems.

Example 3–11 shows a potential hardware problem involving a qna device. In little more than an hour, the qna required three restarts, as indicated by the qna restart messages ❶, ❷, and ❸.

Example 3–11 uerf Example

```
# uerf -R
   uerf version 3.1-003 (113)

***** ENTRY      1. *****
----- EVENT INFORMATION -----
EVENT CLASS                OPERATIONAL EVENT
OS EVENT TYPE              250.  ASCII MSG
SEQUENCE NUMBER           85.
OPERATING SYSTEM          ULTRIX 32
OCCURRED/LOGGED ON        Wed May 30 10:44:34 1990 EDT
OCCURRED ON SYSTEM        beacon.nac.d
SYSTEM ID                  x08000000
SYSTYPE REG.              x01010000

                           FIRMWARE REV = 1.
PROCESSOR TYPE            KA630
MESSAGE                   qerestart: restarted qe0 76 ❶
***** ENTRY      2. *****
----- EVENT INFORMATION -----
EVENT CLASS                OPERATIONAL EVENT
OS EVENT TYPE              250.  ASCII MSG
SEQUENCE NUMBER           84.
OPERATING SYSTEM          ULTRIX 32
OCCURRED/LOGGED ON        Wed May 30 10:04:44 1990 EDT
OCCURRED ON SYSTEM        beacon.nac.d
SYSTEM ID                  x08000000
SYSTYPE REG.              x01010000

                           FIRMWARE REV = 1.
PROCESSOR TYPE            KA630
MESSAGE                   qerestart: restarted qe0 75 ❷
***** ENTRY      3. *****
----- EVENT INFORMATION -----
```

(continued on next page)

Network Management and Troubleshooting Tools

uerf command

Example 3–11 (Cont.) uerf Example

```
EVENT CLASS                                OPERATIONAL EVENT
OS EVENT TYPE                             250.          ASCII MSG
SEQUENCE NUMBER                           83.
OPERATING SYSTEM                          ULTRIX 32
OCCURRED/LOGGED ON                        Wed May 30 09:35:14 1990 EDT
OCCURRED ON SYSTEM                        beacon.nac.d
SYSTEM ID                                  x08000000
SYSTYPE REG.                              x01010000

PROCESSOR TYPE                             KA630
MESSAGE                                    qerestart: restarted qe0 74 ③
#
```

Resources for Troubleshooting

This chapter provides information on the following resources and procedures that are used frequently in troubleshooting network problems:

- ULTRIX and VMS log files
- Routing path trace procedures
- NCP loopback tests
- Reachability tests for TCP/IP hosts
- NCP counters
- ULTRIX counters
- DECnet events

Log Files

This section describes various types of log files that enable you to collect information you can use in troubleshooting.

The log files for VMS systems provide information on system events, user requests, network operations, hardware error messages, and accounting statistics. The log file for ULTRIX systems provides information such as hardware or software error messages, informational messages, emergency messages such as hard disk failures, warnings about abnormal conditions, and debugging information.

VMS OPCOM and the Operator Log File

The operator communication manager (OPCOM) on VMS systems sends information on system events and user requests to the operator terminal and to the operator log file (SYSS\$MANAGER:OPERATOR.LOG). OPCOM can provide information on the events preceding a network problem, as well as information that can help you anticipate and prevent hardware and software failures.

To start OPCOM, enter the following command from an account that has OPER privileges:

```
$ @SYS$SYSTEM:STARTUP.COM OPCOM
```

To enable event logging using OPCOM, run NCP and enter the following commands:

```
NCP> SET LOGGING MONITOR KNOWN EVENTS
NCP> DEFINE LOGGING MONITOR KNOWN EVENTS
NCP> SET LOGGING MONITOR STATE ON
NCP> DEFINE LOGGING MONITOR STATE ON
```

Resources for Troubleshooting Log Files

VMS Netserver Log File

Netserver log files on VMS systems often contain very useful information in addition to that displayed in DECnet error messages. This additional information, obtained from the appropriate node, can help you isolate and resolve DECnet problems.

Netserver log files are located in the account that initiates the login sequence. Depending on how the user tried to access a remote node, the netserver log file can be located in one of the following:

- Default nonprivileged DECnet account
- Object account
- Proxy account
- Any account accessed using explicit access control information

Note

Netserver log files are purged often and are kept open by server processes, so be sure that you are using the correct netserver log file for the problem you are trying to solve.

VMS Error Log File

VMS systems automatically write error messages to the latest version of the SYSSERRORLOG:ERRLOG.SYS file. This file logs all hardware errors for the VMS system, and is useful in diagnosing network problems because it documents the errors and events preceding a failure. The error log file is primarily a Field Service tool.

You can display information from the VMS error log file using the DCL command, SHOW ERROR. To generate a complete report of the errors in the VMS error log file, including a history and a detailed description of the errors, do the following:

1. Make sure that you have SYSPRV privilege. This privilege is required to access the error log.
2. Set default to SYSSERRORLOG.
3. Display the error log directory to see which error log file you want to analyze.
4. Enter the following command to generate a full report of the current error log file:

```
$ ANALYZE/ERROR_LOG/OUTPUT=ERRORS.LIS
```

VMS Accounting Log File

The accounting facility collects statistics on the use of system resources in the accounting log file, SYSSMANAGER:ACCOUNTNG.DAT. The accounting log file contains information on everything for which accounting is currently enabled on the system. Use this information to track job-specific problems, such as those involving NETACP and REMACP.

To use the accounting log file, you must have read access to the file.

ULTRIX Error Log Files

The uerf command in ULTRIX runs the ULTRIX error formatting program and displays the contents of system events recorded in the error log file, /usr/adm/syserr/syserr.hostname. You cannot read the error file without using uerf first to decode it. The error log file includes information about the system hardware and the software kernel, as well as information about system status, startup, and diagnostics.

The syslog daemon in ULTRIX records system messages into a set of files. These messages indicate a variety of system conditions, including error conditions and warning messages, and notices to users.

For more information about uerf and syslog, see Chapter 3.

Routing Path Trace Procedures

Tracing the routing path between two nodes helps you determine the point of failure between the nodes. Performing routing path trace procedures is a manual way of generating network topology information. By understanding what nodes are between the local and remote nodes, you can begin to isolate the problem to the path between the nodes.

In tracing a routing path, you make connections systematically across the network from the local node to each of the nodes along the path to the remote node. If no problems exist along the path, you eventually reach the remote node. However, a node along the path may be unreachable. The unreachable node is the point of failure between the local and the remote node, and is where you need to focus your troubleshooting efforts.

This section provides information on the NCP path trace procedure for DECnet networks, and the netstat trace procedure for TCP/IP networks. For TCP/IP networks, you can also use an unsupported tool called traceroute. See Chapter 3 for more information about traceroute.

Resources for Troubleshooting Routing Path Trace Procedures

NCP Routing Path Trace Procedure

Before you begin to trace a path with NCP, make sure that the local node is running properly, and the remote node is properly defined in the local node's configuration database. If the problem persists after you check the local node and its configuration database, run NCP and do the following to trace the path:

1. Enter the following command to display the name of the next (or adjacent) node on the path from the local node to the remote node:

```
NCP> SHOW NODE remote-node-id
```

2. Use the NCP TELL command to direct the next node (displayed with the previous command) to display the node name for the next node on the path to the remote node.

Use the remote node's node address to ensure that you are tracing the correct node through all the nodes between the local and remote nodes:

```
NCP> TELL next-node-id SHOW NODE remote-node-address
```

3. Repeat step 2 until you reach the remote node, or until NCP displays an error message.
 - If you can reach the remote node, the path between the local and remote node is intact. Any previous failures to reach the remote node may have been caused by a transient or intermittent error.
 - If you can reach all the nodes between the local and remote node, but cannot reach the remote node, the problem is on the remote node.
 - If you cannot complete a connection from one node on the path to the next node on the path, the problem is a circuit failure between the two nodes on the path, or a problem on the next node.

Figure 4–1 shows the routing path trace procedure in flowchart form.

Figure 4–1 NCP Routing Path Trace Procedure

netstat Routing Path Trace Procedure

The netstat command in ULTRIX enables you to obtain routing path information when tools such as traceroute are not available. To get routing path information for a remote host, do the following:

1. Display the local host's routing tables to find the IP router that the local host uses to reach the destination network.
2. Display the routing tables on the local host's IP router.
3. Continue to display the routing tables for each IP router on the path until you reach the destination network and host.

Example 4–1 shows partial routing tables for a host called host5, and an IP router called router1. Assume that host5 needs to send datagrams to a host on network 12.11. For any host not directly connected to host5, the routing tables show that host5 uses a default route (through router1) ❶.

To find the next IP router on the path to network 12.11, use the netstat -r command on router1 to display router1's routing tables

Router1 has multiple interfaces (dmv0, dmv1, dmv2, qe0) to different networks. Router1 forwards the datagrams from host5 through one of these interfaces, depending upon the destination host's address. In this case, router1 sends the datagrams through interface dmv1 to the IP router, ssq2.corp.com❷, to reach network 12. The route to network 12 allows the datagram to reach network 12.11, because network 12.11 is a subnet of network 12.

Resources for Troubleshooting Routing Path Trace Procedures

Example 4–1 netstat Routing Path Trace

```
% netstat -r

Routing tables
Destination      Gateway          Flags    Refcnt  Use      Interface
localhost        localhost        UH       5        11483    lo0
default          router1 ①        G        0         0        ln0
16.20.32         host5           U        8        24255    ln0

% netstat -r

Routing tables
Destination      Gateway          Flags    Refcnt  Use      Interface
router12zso     zsogw           UGH      0         0        dmv2
ucs2ssq.corp.com ssq1.corp.com  UGH      0        38776    dmv0
12              ssq2.corp.com  ②      UG       42       16337    dmv1
localhost       localhost        UH      10       194101    lo0
18.45.224       ssq1.corp.com  UG       0         0        dmv0
default         ssq1.corp.com  UG       0       5981933    dmv0
16.20.32        router1         U        8       2427750    ge0
```

NCP Loopback Tests

Loopback tests apply to problems at the node, LAN, and WAN levels. You can use loopback tests to isolate the source of a problem to the node or component (or both) causing a problem on the circuit path. In loopback testing, a **component** can be the cable, modem, hardware, software, or circuit. Loopback tests verify a circuit or logical link's operational status, with each type of loopback test verifying a different portion of the circuit or logical link. As such, loopback tests show you the components that are functioning properly, and where the circuit path or node fails.

Loopback tests fall into two broad categories: node-level and circuit-level. Use node-level tests to evaluate the operation of logical links, routing, and other network-related software. Use circuit-level tests to evaluate the operation of circuits, or in terms of the DECnet layers—the physical and data link layers.

Loopback tests exercise hardware and software by repeatedly sending data through network components and returning the data to its source. By changing the loopback point, you can locate problems precisely. Figure 4–2 shows the path from node Q to node R over a common carrier, and illustrates the various points (designated as L_n) to which you can loop data to determine the point of failure.

Figure 4–2 Components and Points of Loopback Testing

You can perform loopback testing in one of the following ways:

- Test from a point on the circuit path out to the remote node until a problem occurs

A failed loopback test to one of the components on the path points to the failed component or node causing the initial problem.

- Test from the remote node back to the local node until a test is successful

The first successful loopback test you perform shows you the extent of the circuit that is working properly, and indicates the component or node that is the point of failure and the cause of the initial problem.

Loopback tests are primarily used to troubleshoot WAN problems, especially node reachability problems, and routing problems, such as partitioned areas. For node and LAN problems, loopback tests help verify logical link capabilities as well as hardware and software functionality. In addition, most modems, circuits, cables, and common carriers have loopback capabilities. Some devices allow you to perform loopback tests using a loopback connector.

Resources for Troubleshooting NCP Loopback Tests

Other uses for loopback tests include the following:

- Stress-testing of paths, circuits, or routers
- Evaluating circuits that are not operating optimally
- Testing router throughput (congestion loss)

By specifying the length and count of a loop, you can define in detail how much data should pass through or to the router for a given link. **Length** is the length (in bytes) of the blocks to be sent during loopback testing, and **count** is the number of blocks to be sent during loopback testing.

NCP Loopback Test Results

In a successful loopback test, the data loops back to its source uncorrupted, and the NCP> prompt is displayed. Sometimes the loopback test may encounter **negative acknowledgments (NAKS)** or other problems, but the test is still successful because it loops the correct count and length of packets.

Loopback tests provide only one source of information you need for troubleshooting a problem. Other significant errors may have been logged which indicate a problem. Be sure to check NCP circuit counters and error log files for unusual errors as well. See the previous sections on log files and NCP Counters for more information on NCP counters.

In a failed loopback test, the data either does not return to its source or returns in a corrupted state. The result is an error message rather than the NCP> prompt. For example, a loopback test may fail because all of the packets did not loop successfully. This may indicate congestion or a circuit problem.

If a loopback test fails, keep in mind what the test was intended to accomplish and where the test failed in terms of node and DNA layer. This information may be useful as you attempt to isolate the problem further. Use the error log files (described previously) and DECnet counters (see NCP Counters) to gather more information about the problem.

Types of Loopback Tests

Loopback tests fall into two broad categories: node and circuit. Within the node and circuit categories, loopback tests are further classified as either hardware or software loopback tests, and as disruptive or nondisruptive, depending on their effect on network operation. The following list items describe the various types of loopback tests:

- **Hardware loopback tests** loop data through specific hardware points using a hardware loopback connector inserted on the end of a cable. Hardware loopback tests loop data back without any modification.
- **Software loopback tests** loop data through various DNA layers using NCP commands. Software loopback tests may amplify or regenerate the signal before looping the data back to the source.
- **Disruptive loopback tests** disconnect the existing circuit before looping the data, either physically with a loopback connector or button, or through software. Disruptive loopback tests disable any services the circuit may have been providing.

- **Nondisruptive loopback tests** do not disconnect the circuit before looping the data, and are particularly helpful in isolating intermittent or transient problems. For example, a user calls with a node unreachable problem, but when you try the same operation, the node is reachable. This type of transient or intermittent problem can be fairly difficult to isolate. However, using loopback tests, you may be able to cause this error to occur again.

As a guideline for deciding whether to use disruptive or nondisruptive loopback tests, keep the following in mind:

- If the circuit is up and providing services, leave the circuit up and perform nondisruptive loopback tests.
- If the circuit is alternating between states (also known as a **bouncing circuit**) and causing significant problems, or if the circuit is down, you can use the disruptive tests to isolate the problem because you need not be concerned about disrupting service that is already interrupted.

Node-Level Tests

Node-level loopback tests examine the logical link capabilities of a node by exchanging test data between DECnet processes on two different nodes or between DECnet processes on the same node. The two types of node-level loopback tests are:

- Loopback tests for logical link operation over unspecified circuits
- Loopback tests for operation over a specified circuit

If the loopback test fails, NCP displays a message that indicates a test failure, specifies the reason for the failure, and provides a count of the messages that were not returned. For a summary of NCP error messages, see the *VMS System Messages and Recovery Procedures Reference Volume*.

Remote Loopback Test

Use the remote loopback test to verify the logical link connection between two nodes. For example, to test the local and remote DECnet software on nodes BOSTON and TRNTO, use the NCP commands shown in Figure 4-3.

Resources for Troubleshooting NCP Loopback Tests

Figure 4–3 Remote Loopback Test

Local and Remote Loopback Tests Using a Loop Node Name

If the remote loopback test fails, then use the LOOP NODE command with a loop node name to test a logical link path over a specified circuit. You can loop test messages in one of the following ways:

- Over a logical link path and circuit within the local node
Use this method first to test the remote routing layer software.
- Between two different nodes with a loop node specified for the circuit to be used

In each case, use the SET NODE command with the CIRCUIT parameter to establish a loop node name. For example, the following command establishes circuit DMC-0 as the circuit over which loop testing will take place:

```
NCP> SET NODE tester CIRCUIT dmc-0
```

The following restrictions apply to loop node tests:

- CIRCUIT is the only valid parameter for loop nodes.
- The circuit specified must be on when performing LOOP NODE tests.
- Each circuit can have only one loop node name at a time.

For example, after you establish TESTER as the loop node name for circuit DMC-0, you must use the CLEAR NODE TESTER CIRCUIT command before assigning another loop node name to DMC-0.

When a logical link connection request is made to the loop node name, all subsequent logical link traffic is directed over the associated circuit. The destination of the traffic is the node address associated with the loop node name. The loop node name is necessary because, under normal operation, DECnet routing software selects which path to use when routing information. The loop node name overrides the routing function so that information can be routed over a specific circuit. To remove the association of the loop node name with a circuit, use the CLEAR NODE CIRCUIT or CLEAR NODE ALL command, as in the following example:

```
NCP> CLEAR NODE tester CIRCUIT
```

Local-to-Remote Testing

Use local-to-remote loopback tests to verify the logical link path over a circuit between the local node and a remote node, as shown in Figure 4–4.

The commands in Figure 4–4 test both local and remote routing layer software operation. The test messages are looped over the loopback circuit. Because the test actually verifies the operation of the routing layer on the remote node, the message might not come back on the circuit over which it was sent.

Figure 4–4 Local-to-Remote Loopback Test Using a Loop Node Name

If the local-to-remote test fails, use a local-to-local loopback test on both the local and remote nodes to verify the logical link path over a specified line, as shown in Figure 4–5.

The example in Figure 4–5 verifies only the local routing layer software. Because the device is set to loopback mode, the test messages are looped over the circuit and back to the local node. A failure of this test indicates a problem on the node on which the test was executed.

Note

The local-to-remote loopback test is a disruptive loopback test.

Also, because of restrictions in the operation of the DMC controller, you must use a block length of fewer than 50 bytes for controller loopback tests.

Resources for Troubleshooting NCP Loopback Tests

Figure 4–5 Local-to-Local Loopback Test Using a Loop Node Name

Local Loopback Test

If the local-to-local loopback tests fail, use a local loopback test to verify the local DECnet software, as shown in Figure 4–6.

This test evaluates the local DECnet software using an internal logical link path. If this test succeeds and the other node-level tests fail, then try the circuit-level tests. A failure of this test indicates a problem on the local node. For example, a failure might indicate the improper setup of the account or file protections for the object or default nonprivileged DECnet account.

Figure 4–6 Local Loopback Test

Circuit-Level Tests

Circuit-level loopback tests examine a DECnet circuit by looping test data through a hardware loopback device on the circuit, either through a modem (or loopback connector) or through a remote node. The tests that use a hardware loopback device are referred to as **controller loopback tests**. The tests that use a loopback connector or a modem are referred to as **circuit loopback tests**. The tests that use the software capabilities of the system are referred to as **software loopback tests**.

You can also perform circuit-level loopback tests for Ethernet circuits.

Note

By default, the SERVICE parameter is set to DISABLED for Ethernet circuits. However, for Ethernet circuit-level loopback testing, an Ethernet circuit must be in the ON state and the SERVICE parameter must be set to ENABLED.

Using Ethernet circuit-level loopback tests, you can specify physical and hardware address parameters. You can also request the help of an assistant physical address or node in completing the test, as follows:

- If you have trouble transmitting a message to a target node, you can request assistance in transmitting the message to the target node.
- If you can transmit messages to the target node, but not receive messages from it, you can request assistance in receiving a message from the target node.
- If you have difficulties both in sending and receiving messages, you can request assistance both in transmitting messages to and receiving messages from the target node.

The following sections describe each type of circuit-level loopback test.

Resources for Troubleshooting NCP Loopback Tests

Software Loopback Test

Use the software loopback test to loop through the circuit-to-circuit service software in the adjacent node and back to the local node, as shown in Figure 4-7. The software loopback test checks whether the circuit is operational up to the remote unit and controller on the adjacent node.

The commands in Figure 4-7 test the circuit DMC-0 up to the adjacent node. If this test fails, try a circuit loopback test to verify that the circuit is functional.

Figure 4-7 Software Loopback Test

Controller Loopback Test

Figure 4-2 in NCP Loopback Tests illustrates the components and points of loopback testing. When you perform controller loopback tests, you perform the test to each point (L_n) along the path, as shown in Figure 4-2, until you determine the point of failure.

Use the controller loopback test to verify whether the circuit up to the controller and the controller itself are functional, as shown in Figure 4-8.

The commands in Figure 4-8 test the circuit up to the controller for physical line DMC-0 connected to the local node by circuit DMC-0.

Note

The controller loopback test is a disruptive loopback test.

Also, because of restrictions in the operation of the DMC controller, you must use a block length of fewer than 50 bytes for controller loopback tests.

Figure 4–8 Controller Loopback Test

Modem Loopback Tests

Modem loopback tests allow you to test from the local node to the modem and back, and from the remote node to the modem and back. These tests help determine whether a problem is between the node and the modem, or elsewhere on the path.

Most modems have local and remote loopback capabilities. You can use local modem loopback tests to isolate problems between the modem and the circuit.

You can also use a loopback connector to perform loopback tests through a modem. A loopback connector loops data through all the hardware in the path of the loopback test.

For example, if you use a loopback connector to test the configuration in Figure 4–2 in NCP Loopback Tests, you use it at points L_3 , L_4 , L_{10} , and L_{11} . You can use the loopback tests to perform both local and remote looping through the loopback connector, or simply remote loopbacks as a function of a modem.

Using Loopback Tests to Check Circuitry

Most point-to-point problems involve the circuitry. The following procedure shows you how to perform loopback tests to determine whether the circuitry is working properly. See Figure 4–2, Components and Points of Loopback Testing, for the configuration used in this example.

1. Perform local loopback testing on node Q at the modem. (This procedure takes the circuit down.)

If the test completes successfully, node Q and the modem are operating properly. Go to step 2.

If the test fails, go to step 3.

2. Perform local loopback testing on node R at the modem.

Resources for Troubleshooting NCP Loopback Tests

If the test completes successfully, node R and the modem are operating properly. The problem is probably the circuit between the modems unless there are hardware setup or verification problems (such as incorrectly specified transmit and receive passwords on the circuit. These kinds of problems are indicated in the operator log file.) Go to step 5.

If the test fails, go to step 3.

3. Perform local loopback testing at the interface between the communications board and the cable to the modem.

If the test completes successfully, the problem is the cabling between the interface and the modem. Go to step 4.

If the test fails, the problem is a hardware failure at the device level. Call Field Service.

4. Perform local loopback testing at the cabling just before modem.

If the test completes successfully, do the modem self-test.

If the test fails, then replace the cabling.

5. Perform remote testing on the circuit, using either a remote loopback with the modem or a remote loopback using a loopback connector.

This test loops data through the circuit and shows if the circuit is working properly. If the test completes successfully, check the verification (transmit and receive passwords) and switch settings at the board level on each node.

If the test fails, have the circuit repaired.

Figure 4-9 shows the procedure in flowchart form.

Figure 4–9 Modem Loopback Tests

Reachability Tests for TCP/IP Networks

The ping command tests the reachability of hosts on a network using ICMP. The ping command is useful for troubleshooting direct and indirect routing problems.

Resources for Troubleshooting Reachability Tests for TCP/IP Networks

In addition to using ping to test low-level functions, you can try to log in to the host in question using commands such as rlogin and telnet to check higher-level problems such as the lack of an rlogin or telnet server.

To diagnose TCP/IP problems, you can use the software loopback interface. On ULTRIX systems, this interface is "ln0," and always has the address 127.0.0.1, and the name "localhost." You can use this interface with the ping command or the rlogin command to determine the network reachability of the local host, for example:

```
# ping localhost
localhost.corp.com is alive
```

For more information about the ping command see Chapter 3.

NCP Counters

Even when a loopback test completes successfully, problems may still exist, especially if the problem is related to DECnet or to hardware errors. You can gather additional troubleshooting information by evaluating NCP counters. NCP counters can provide vital troubleshooting information in addition to the information obtained from loopback testing.

When you use counters, you need to know the length of time counter statistics have been accumulating. So, before working with counter information, always check to see how much time has elapsed for a counter. By doing this, you can see how much data the test passes. The amount of data passed should be approximately as follows:

data passed \geq (message length \times messages looped)
or
data passed \geq (length \times count)

Length is the length (in bytes) of the blocks to be sent during loopback testing, and **count** is the number of blocks to be sent during loopback testing.

Note

NCP displays the greater-than symbol (>) when a counter has reached its maximum and cannot record any further changes. If a counter displays a greater-than symbol (>), the counter information is not very useful, because you cannot tell how long the counter information has been accumulating or how much the counter has incremented. For this reason, it is a good idea to zero the counters automatically and periodically. To avoid problems when troubleshooting, always check the counters before and after the test, or zero the counters before performing tests.

In troubleshooting, the following counters can provide critical information for problems dealing with circuits, lines, and nodes:

- Seconds since last zeroed
- Terminating packets received
- Originating packets sent
- Transit packets sent and received
- Bytes received and sent

- Data blocks received and sent

The following sections describe particular counters as they apply to node, LAN, and WAN problems.

Errors Applicable to Node Problems

Errors of special interest when troubleshooting node-level problems include the following:

- Buffer unavailable
- Packet format error
- Received connect resource errors
- Response timeouts

Errors Applicable to LAN Problems

In LAN environments, the following errors are all Ethernet problems that indicate the problem is LAN and hardware related:

- Unrecognized frame destination errors
- Collision detect check failures
- Send and receive failures
- Circuit down failures

Errors Applicable to WAN Problems

Errors of special interest when troubleshooting WAN problems include the following:

- Circuit down
- Data errors inbound and outbound
- Local and remote reply timeouts
- Transit congestion loss (routers only)

These errors are routing errors showing that a circuit or router is too busy to properly handle the network traffic.

Resources for Troubleshooting NCP Counters

In WAN environments, errors point to the local or the remote node as the point of failure. The following errors may occur on the local node:

- Data errors outbound
- Local buffer errors
- Local reply timeouts

The following errors point to problems on the remote node:

- Data errors inbound
- Remote buffer errors
- Remote reply timeouts

If the errors point to both the local and the remote nodes, then the problem is probably between the two nodes, for example, the circuit, modems, or terminating equipment. Errors that point only to the local or only to the remote node may be caused by circuit problems on the respective side.

Errors Applicable to Node, LAN, and WAN Problems

"Buffer unavailable" errors affect nodes, LANs, or WANs, and show that the system is busy and needs to be tuned by increasing certain parameters. For "user buffer unavailable" errors, increase the receive buffers on the circuit in question.

Formulas for Understanding Counters

NCP counters provide a large amount of raw data. Sometimes that information is useful in and of itself. Sometimes, however, you may need to perform some calculations using the counter information to better understand the operation of your network. This section provides formulas for calculating the following:

- Packet rate
- Circuit quality
- Transit congestion loss
- Ethernet line statistics
- Retransmissions
- Routing overhead

Packet Rate

Packet rate is the number of packets transmitted per second, and generally applies to circuits. Typically, bridges, routers, and gateways are limited by the packet rate internally and the bandwidth externally. By knowing the packet rate of each network component, you can determine where problems lie and make plans for upgrading hardware.

The formula for packet rate is as follows:

$$\text{packet rate} = (\text{packets sent} + \text{packets received}) / \text{seconds since last zeroed}$$

The formulas for determining packets sent and packets received are as follows:

$$\text{packets sent} = (\text{originating packets sent} + \text{transit packets sent})$$

$$\text{packets received} = (\text{terminating packets received} + \text{transit packets received})$$

You can also determine the total number of blocks that have been sent or received on a particular line. The block rate may be higher than the packet rate, as a result of routing overhead and HELLO messages.

Because most circuits are rated in bits per second (b/s), use the following formula to determine the data throughput on each circuit.

To obtain the total bits per second, including overhead, add (8 bytes \times blocks sent/received) to the calculation. When you measure traffic on asynchronous circuits, reduce the circuit saturation speed by 20 percent to account for the start and stop bits.

$$\text{b/s out} = (\text{bytes sent} \times 8) / \text{seconds since last zeroed}$$

$$\text{b/s in} = (\text{bytes recd} \times 8) / \text{seconds since last zeroed}$$

Circuit Quality

Circuit quality is the percentage of data transmitted that reaches its destination intact, and is based on the throughput through the circuit, and data errors and timeouts recorded. The circuit quality formulas provide a way to compare the performance of circuits against a standard. You can use the circuit quality formula to evaluate both circuits and lines.

Use the following formulas to calculate the circuit quality as a percentage value:

$$\text{circuit quality in} = 100 \times \text{blocks sent} / (\text{blocks sent} + \text{errors out} + \text{local timeouts})$$

$$\text{circuit quality out} = 100 \times \text{blocks recd} / (\text{blocks recd} + \text{errors in} + \text{remote timeouts})$$

Other ways to look at errors on a circuit are on the basis of per hour or per megabyte (mbyte) of data. The following formulas show how to calculate errors per hour and errors per megabyte of data:

$$\text{errors per hour inbound} = (\text{errors in} + \text{remote timeouts}) / \text{hours since last zeroed}$$

$$\text{errors per hour outbound} = (\text{errors out} + \text{local timeouts}) / \text{hours since last zeroed}$$

Resources for Troubleshooting NCP Counters

errors per mbyte outbound = $10E5 \times (\text{errors out} + \text{local timeouts}) / \text{bytes sent}$

errors per mbyte inbound = $10E5 \times (\text{errors in} + \text{remote timeouts}) / \text{bytes recd}$

Transit Congestion Loss

Transit congestion loss is the state where circuits or routers or both receive too many packets to be processed at one time. When this occurs, the packets are discarded, and the transport layer of the system that originated the packets retransmits the packets.

Transit congestion loss acts as a pressure relief valve for overloaded circuits or routers. As a guideline, one to two percent congestion loss is considered normal on networks that attempt to maximize the throughput of their data circuits.

The following formula shows how to determine the percentage of loss:

transit congestion loss = $100 \times \text{transit congestion} / (\text{packets sent} + 1)$

Ethernet Line Statistics

The Ethernet line statistics that are important for troubleshooting include the percent of deferred packets and the percent of collisions. These statistics are the result of calculations involving counters that keep track of the way packets gain access to the cable.

Packets waiting to be put on the cable must wait for channel availability. If there is a delay at this stage, the initially deferred counter is incremented. If the packet is put on the cable and experiences a collision, the back off and retransmit sequence starts and the single collision counter is incremented. If more than 16 collisions occur before a packet can be put on the cable, the send failure counter is incremented.

In a typical network, collisions remain under 10 percent of the total traffic sent over any Ethernet controller. Use the following formula to calculate Ethernet line statistics (based on data received from the NCP command, SHOW LINE COUNTERS):

percent deferred = $100 \times (\text{initially deferred} + \text{single collision} + \text{multiple collision}) / (\text{data blocks sent} + \text{multicast blocks sent} + 1)$

percent collisions = $100 \times (\text{single collision} + \text{multiple collision}) / (\text{data blocks sent} + \text{multicast blocks sent} + 1)$

Keep in mind that the traffic from one node does not reflect the traffic on the entire Ethernet. When measuring b/s throughput, packet rates, or collision rates from any one system, you see only the effect of that system on the Ethernet. Total Ethernet usage, packet rate, and collision rate can be determined only by Ethernet monitoring tools and applications.

Retransmissions

A certain number of retransmissions between any two systems is normal. The rate of retransmissions depends on the size and length of the network. High numbers of retransmissions, however, indicate problems somewhere along the path. Use the following formula to calculate the retransmission rate (based on data received from the NCP command, SHOW NODE node-id COUNT):

$$\text{percent transport retransmissions} = 100 \times (\text{response timeouts} / \text{messages sent})$$

Routing Overhead

Routing overhead is the amount of traffic generated to maintain the adaptive routing tables. The formula for calculating routing overhead measures the amount of traffic across a data link that comes from above the network application layer, and how much is generated to maintain the adaptive routing tables. This value can be calculated on both point-to-point and Ethernet circuits. Use the following formula (based on data received from the NCP command, SHOW CIRCUIT circuit-id COUNT):

$$\text{overhead percent} = 100 \times (\text{total blocks—total packets}) / (\text{total blocks})$$

ULTRIX Counters

A variety of SNMP-based tools provide counter information, and calculations based on those counters for ULTRIX systems. See the documentation that accompanies the tool used at your site for more information on ULTRIX counters.

DECnet Events

The DECnet event logging facility (EVL) monitors significant network events such as circuit failures or lost packets, on a continuous basis. An **event** is a network or system-specific occurrence for which the logging component maintains a record. This information can help in isolating the source of network problems.

The following is a partial list of significant events logged:

- Circuit and node counter activity
- Changes in circuit, line, and node states
- Service requests (when a circuit or line is put in an automatic service state)
- Passive loopback (when the executor is looping back test messages)
- Routing performance and error counters (circuit, line, node, and data packet transmission)
- Data transmission performance and error counters (when errors in transmission occur)
- Lost event reporting (when some number of events are not logged)

Resources for Troubleshooting DECnet Events

To enable event logging, take the following steps from an account that has OPER privileges:

1. Enter the following command to check whether OPCOM is running:

```
$ SHOW SYSTEM
```

2. If OPCOM is not one of the process names shown, enter the following command to start it:

```
$ @SYS$SYSTEM:STARTUP.COM OPCOM
```

3. Run NCP, and enter the following commands to enable OPCOM to perform normal event logging:

```
NCP> SET LOGGING MONITOR KNOWN EVENTS  
NCP> DEFINE LOGGING MONITOR KNOWN EVENTS  
NCP> SET LOGGING MONITOR STATE ON  
NCP> DEFINE LOGGING MONITOR STATE ON
```

4. Enter the following DCL command to enable the display of network events at your terminal:

```
$ REPLY/ENABLE=NETWORK
```

Network Troubleshooting Procedures

This chapter discusses common network problems. It presents the problems according to the symptom each problem displays. A symptom can be a displayed message (such as the DECnet event message, "Aborted Service Request"), or a description of the problem situation (such as, "LAN Bridge Cannot Downline Load").

Following each symptom is an explanation of why the problem occurs, a brief description of the troubleshooting strategy for the problem, a troubleshooting procedure that gives step-by-step instructions for solving the problem, and general recommendations, if any.

The step-by-step procedures do not present a complete methodology for solving the problems. Instead, they provide the *most likely* solutions for the problems—the solutions to try first.

For problems not addressed in this chapter, see Chapter 2, Network Troubleshooting Methodology, for other problem-solving approaches.

Organization

The network problems discussed in this chapter are organized alphabetically according to the symptom message or problem description.

The explanation section for each problem gives the reasons for the problem and describes the extent of the problem. Table 5-1 categorizes the problems according to the extent of their disturbance on the network, as follows:

- ULTRIX host problems
- VMS node problems
- LAN problems
- WAN problems
- Cross-category problems (two or more of the preceding problems)

Table 5-1 also includes the protocol involved in each problem.

Network Troubleshooting Procedures Organization

Table 5–1 Network Problems and Extent of Disturbance on the Network

Extent of Problem	Protocol	Symptom or Problem Description
ULTRIX Host	DECnet	Connection Failed, Access Control Rejected
ULTRIX Host	DECnet	Connection Failed, Unrecognized Object
ULTRIX Host	IP	Login Incorrect
ULTRIX Host	IP	Permission Denied
ULTRIX Host	IP	Unknown Host
VMS Node	DECnet	Device Not Mounted
VMS Node	DECnet	Insufficient Resources At Remote Node
VMS Node	DECnet	Invalid Parameter Value
VMS Node	DECnet	Line Synchronization Lost
VMS Node	DECnet	Login Information Invalid
VMS Node	DECnet	Network Object Unknown
VMS Node	DECnet	Network Partner Exited
VMS Node	DECnet	Node Out Of Range Packet Loss
VMS Node	DECnet	Partial Routing Update Loss
VMS Node	DECnet	Verification Reject
LAN	DECnet/MOP	Aborted Service Request
LAN	DECnet	Adjacency Rejected, Adjacency Up
LAN	Generic	Babbling Device
LAN	Generic	Broadcast Storm
LAN	DECnet/MOP	LAN Bridge Cannot Downline Load
LAN	Generic	LAN Segment Communication Problem
LAN	LAT	LAT Port Hung
LAN	LAT	LAT Print Queue Problems
LAN	LAT	Terminal Server Connection Failures
WAN	DECnet	Asynchronous DECnet Problems
WAN	IP	Network is Unreachable
WAN	DECnet	Partitioned Area
Cross-category	DECnet	Circuit State Problems
Cross-category	IP	Connection Timed Out
Cross-category	Generic	Dialup Problems
Cross-category	IP	Host is Unreachable
Cross-category	DECnet	Remote Node is not Currently Reachable

Troubleshooting Notes

Keep the following in mind as you begin troubleshooting network problems:

For VMS Systems

- Using privileged accounts
Many of the procedures in this chapter require the use of an account with system management level privileges. Many procedures also assume that you have access to accounts with these privileges on all nodes in your network. For procedures requiring use of a privileged account on a node to which you do not have access, ask the system manager of that node to perform the action.
- Using the SET and DEFINE commands in NCP
Many procedures call for setting NCP parameters with the SET command, which modifies only the volatile database. After you verify that the solution for a problem works, be sure to use the DEFINE command to modify the permanent database.
- Modifying passwords
Some procedures call for correcting mismatches between passwords specified in the SYSUAF file and NCP databases. Before you make any changes, be sure that you have the authority to make the modifications. If you do not have the authority to do so, refer the required change to the appropriate system or network manager.

For ULTRIX Systems

- Modifying passwords
Some procedures call for correcting mismatches between passwords specified in the /etc/passwd file and NCP object databases. Before you make any changes, be sure that you have the authority to make the modifications. If you do not have the authority to do so, refer the required change to the appropriate system or network manager.

Aborted Service Request

Aborted Service Request

Symptoms

With the local Ethernet circuit service state enabled, the system displays one of the following DECnet event messages:

```
%%%%%%%%%% OPCOM 5-OCT-1988 13:48:07.73 %%%%%%%%%%%  
Message from user DECNET on NODE1  
DECnet event 0.7, aborted service request  
From node x.xxx (NODE1), 5-OCT-1988 13:48:07.73  
Circuit UNA-1, Line open error, . . .
```

```
%%%%%%%%%% OPCOM 5-OCT-1988 13:48:07.73 %%%%%%%%%%%  
Message from user DECNET on NODE1  
DECnet event 0.7, aborted service request  
From node x.xxx (NODE1), 5-OCT-1988 13:48:07.73  
Circuit UNA-1, Receive timeout, . . .
```

Explanation

These messages indicate a LAN problem involving DECnet and the MOP protocol. The problem can affect any load host system. A **load host system** is any system that provides downline loading and upline dumping for other systems.

Generally, this symptom results from a node requesting a service from an adjacent node. However, a problem prevents the request from being processed at the adjacent node.

- A. In the case of the *line open error* message, the Network Management Listener (NML) on the adjacent node receives a maintenance operation protocol (MOP) message, but is unable to acquire control of the line. NML on the adjacent node scans the node database, but is unable to locate a matching hardware address for the requesting node.

A load host system can be set up to allow loading for some systems and not for others. If a node is receiving requests when it is not set up to load other systems (for example, when it does not have the appropriate files), you can disable loading or just ignore any inappropriate load requests.

Finally, if the adjacent node is *not* intended to receive MOP requests, but is still receiving requests, the adjacent node may be set up improperly to prevent these requests. You can disable loading, or ignore the aborted service request messages.

Note

Before you take any action on this problem, determine the intended use of the adjacent node.

Aborted Service Request

If the adjacent node *is* intended to load other nodes, the line open error message can be caused by any of the following:

- Incorrect information in the adjacent node's volatile node database regarding the node that is requesting the operation
Some devices, such as the DECserver 100s and DECserver 200s do not have to be defined in the NCP database on the load host. However, other devices, such as DECserver 500s, DECSAs, DECrouters, and MicroVAXes must be defined.
Some of the information that is commonly missing from or incorrect in the volatile node database includes the following:
 - Hardware address
 - Ethernet address
 - Load file
- Improper protection on the load file, or a nonexistent load file
- Problems with the load image
For example, the load image may not exist, or it may not be readable. If secondary and tertiary files are required, they may not exist or be readable.

B. In the case of the *receive timeout* message, one of the following is occurring:

- The line message receive timer expires before the request can be received from the adjacent node.
- If another node is set up as a load host system for the same server, and the other node services the request first, the remaining load host system receives the receive timeout message. This may not be a problem.

These situations can be caused by any of the following:

- The service timer on the local node is too short.
- The line error level on the load host is too high for any message to get through.
- A hardware problem exists with the node being loaded.
For example, there might be a problem with a QNA.
- A problem exists with the path to the node being loaded.

Troubleshooting Strategy

- A. To resolve the line open error message problem, you may need to perform some or all of the following steps:
- If the node receiving the load requests is not intended to load other nodes, disable loading on the node receiving the requests.
 - Check the local node's node definitions for the requesting node, and update them, if necessary.
 - Check the load file's existence, file protections, and logical name definitions.

Aborted Service Request

- If secondary and tertiary load files are required, check their existence, file protections, and NCP definitions.
- B. To resolve the receive timeout message problem, you may need to perform some or all of the following steps:
 - Perform hardware tests on the node being loaded.
 - Check for problems on the path.
 - Check the service timer and increase it, if necessary.
 - Check the line error level on the load host.

Troubleshooting Procedure

- A. **To resolve the line open error message problem, do the following:**

Note

Step 1 turns the circuit off and causes users to lose connections.

1. If the local node is *not* intended to load other nodes, use the following commands to disable loading on the local node:

```
NCP> DEFINE CIRCUIT circuit-id SERVICE DISABLED
NCP> SET CIRCUIT circuit-id STATE OFF-
NCP> SET CIRCUIT circuit-id ALL
```
2. Use the following NCP command on the node receiving the request to make sure that its volatile node database contains correct information for the requesting node:

```
NCP> SHOW NODE remote-node-id CHARACTERISTICS
```
3. If the requesting node is not defined in the receiving node's volatile node database, or if it is incorrectly defined, use the following command to define it:

```
NCP> SET NODE nodename ADDRESS address SERVICE-
_NCP> CIRCUIT ethernet-device HARDWARE ADDRESS-
_NCP> ethernet-address LOAD FILE file specification
```
4. Use the following DCL command to check whether the load file exists and, if so, whether it has WORLD:READ access specified:

```
$ DIRECTORY/PROTECTION filename.ext
```
5. If the file protection on the load file does not include world read privilege, use the following command to set W:R protection on the file:

```
$ SET PROTECTION filename.ext/PROTECTION=(W:R)
```
6. If the requesting node is a DECserver, make sure the logical name definition, MOM\$LOAD, is defined to point to the directory where the load file is located:

```
$ DEFINE/SYSTEM/EXECUTOR MOM$LOAD directory_specification
```

Aborted Service Request

7. If the requesting node requires secondary or tertiary load files, run NCP and use the following command to check whether the files are defined in the node database:

```
NCP> SHOW NODE node-id CHARACTERISTICS
```

8. If the files are not defined in the node database, use the following command to define them:

```
NCP> SET NODE nodename SECONDARY LOADER file specification,-  
_NCP> TERTIARY LOADER file specification
```

9. If secondary or tertiary load files are required, make sure the file protection on these files includes WORLD:READ access.

B. To resolve the receive timeout message problem, do the following:

1. Use appropriate hardware tests to determine if a problem exists with the node being loaded.
2. If a problem exists with the path to the node being loaded see Chapter 4 for information on resolving path problems.

The problem may be a LAN segment problem such as one caused by a bridge or repeater. You may need to refer to other LAN problems described in this chapter, such as, "LAN Segment Communication Problem", and "Babbling Device."

3. Check the current value of the service timer on the line using the following NCP command:

```
NCP> SHOW LINE line-id CHARACTERISTICS
```

4. If the value for the service timer is too low, increase it using the following NCP command:

```
NCP> SET LINE line-id SERVICE TIMER milliseconds
```

5. Use the following commands to display the load host's line error level:

```
NCP> SHOW LINE COUNTERS  
NCP> SHOW CIRCUIT COUNTERS
```

High error rates can indicate that the load host is unable to load devices. Perform loopback tests on the devices to see if the devices have a problem.

Adjacency Rejected/Adjacency Up

Adjacency Rejected/Adjacency Up

Symptoms

A node repeatedly displays the following DECnet event messages:

```
%%%%%%%%%% OPCOM 27-JUN-1988 09:32:32.98 %%%%%%%%%%%  
Message from user DECNET on NODE1  
DECnet event 4.16, adjacency rejected  
From node x.xxx (NODE1), 27-JUN-1988 09:32:32.82  
Circuit BNT-0, Adjacent node = y.yyy (NODE2)  
  
%%%%%%%%%% OPCOM 27-JUN-1988 09:32:32.93 %%%%%%%%%%%  
Message from user DECNET on NODE1  
DECnet event 4.15, adjacency up  
From node x.xxx (NODE1), 27-JUN-1988 09:32:32.83  
Circuit BNT-0, Adjacent node = z.zzz (NODE3)
```

This symptom can be accompanied by application error messages such as "Path to network node lost." This problem is also known as a **bouncing circuit**.

Explanation

These messages indicate a LAN problem involving the DECnet routing layer. The problem results from conflicts between the designated routing node on the LAN and another routing node on the LAN.

A LAN can consist of multiple DECnet areas, and each DECnet area on a LAN has a designated routing node. Each routing node broadcasts its designated router status across the LAN. However, one of the routing nodes (usually the one that is not the designated routing node) has a hardware problem that causes it to have inaccurate routing information.

This problem tends to occur on QNA systems used as routing nodes. With these systems, a receiver lockup can occur, causing the QNA system to send but not receive routing information. As a result, the QNA system does not get the correct routing information.

Troubleshooting Strategy

To solve this problem, determine the routing node or nodes that are causing the problem, check that the cables are secure, and check for hardware problems on the routing nodes.

Troubleshooting Procedure

1. Look at the OPCOM messages on a system console on one of the end nodes in the LAN. The alternating messages "Adjacency rejected" and "Adjacency up" include the name of the routing node causing the problem.
2. Make sure that the cable connections for the problem routing node are secure.
3. Run NCP on the problem routing node, and use the following command to check whether the problem routing node has a hardware problem:

```
NCP> SHOW KNOWN CIRCUITS
```

Adjacency Rejected/Adjacency Up

If NCP does not list adjacencies, then the board in this system is probably faulty. In this case, shut down the network on this system using the following command:

```
NCP> SET EXECUTOR STATE OFF
```

4. Run diagnostics on the Ethernet interface board to determine the cause of the hardware failure.

Recommendations

Be sure that you define a designated routing node for your LAN. If possible, the routing node should only process routing requests. This is because routing can create a high traffic load on the designated routing node. If the designated routing node is also used for other user activities, response time may suffer.

Some additional guidelines for routing node setup include the following:

- Define the designated routing node's routing priority as 127.
- Define a backup designated router for each LAN.
- Define the backup routing node's routing priority as 126.
- Give the designated routing node the highest DECnet address in its DECnet area.

You do this because designated router status in a LAN defaults to the node with the highest DECnet address in the LAN if two nodes have the same routing priority.

Asynchronous DECnet Problems

Asynchronous DECnet Problems

Symptoms

An asynchronous DECnet connection cannot be created. The circuit is either never established, or it is established, but is in an on-starting or on-synchronizing state.

Explanation

This symptom indicates a WAN problem because it involves a point-to-point circuit. Both static and dynamic asynchronous DECnet can be used in a dedicated circuit, hardwired, or dialup environment. This example assumes that the circuit is a dialup circuit.

Two types of asynchronous DECnet exist: static and dynamic. Static asynchronous DECnet is a connection that allows only DECnet traffic. Dynamic asynchronous DECnet allows both DECnet or simple asynchronous connectivity to occur as needed.

Troubleshooting Strategy

Most asynchronous DECnet problems are the result of improper setup or installation of various devices. Solving this problem requires that you address three potential problem areas:

- A. The system software, which may require that you install additional software and set up system parameters properly

For example, for both dynamic and static asynchronous DECnet, the following need to be properly set up on a VMS system:

- The asynchronous DDCMP driver needs to be loaded into memory.
- The terminal characteristics need to be properly set.

- B. The DECnet software, which may require that you adjust certain network definitions

- C. The communications link, which may require that you address dialup problems

In troubleshooting this problem, use dynamic asynchronous DECnet if possible. Dynamic asynchronous DECnet allows you to establish a connection to a remote node, which shows that data can flow between the two points. This proves that the communication link is operating properly and the system setup is correct.

If you use hardwired or dedicated circuits, use static asynchronous DECnet, and treat this as a circuit problem, solving the problem with loopback tests at the appropriate location. See NCP Loopback Tests and "Circuit State Problems," in this chapter for more information.

Troubleshooting Procedure

A. Check the system software, using the following steps:

1. For both dynamic and static asynchronous DECnet, do the following to install the asynchronous DDCMP driver into memory:

- a. Enable privileges on your process using the following command:

```
$ SET PROCESS/PRIVILEGES=ALL
```

- b. Use the following command to check whether the driver is installed:

- For static asynchronous DECnet, check the system response for the NOA0 device ❶ as shown in Example 5–1.
- For dynamic asynchronous DECnet, check the system response for the VTA0 device ❷ as shown in Example 5–1.

```
$ MCR SYSGEN
SYSGEN> SHOW/DEVICE
```

Example 5–1 SYSGEN SHOW/DEVICE Display

```

__Driver__  __Start__  __End__  __Dev__  __DDB__  __CRB__  __IDB__  __Unit__  __UCB__
NODRIVER    80271680  80274E60
              NOA❶  803D5C40  803BF1C0  803DF480
                                     0  8026E930
LTDRIIVER   80269600  8026DA60
              LTA  803DBFA0  803C0710  803BFCC0
                                     0  80265050
                                     1  80265320
                                     4  8026FBE0
CTDRIVER    80262C90  80264A50
              RTB  803D1C20  803BBB70  803D1E60
                                     0  80243F00
RTTDRIVER   80262150  80262C90
              EIGHT 803D2F40  803BD0C0  803CFEE0
                                     0  80243230
                                     0  8026E930
.
.
.
RTTDRIVER   80255860  8025B739
              VTA❷  803D0EA0  803C16E0  803CF520
                                     0  80227C00
.
.
.
SYSGEN>
```

- c. If SYSGEN does not display the NOA0 device, use the following command to install it:

```
SYSGEN> CONNECT NOA0/NOADAPTER
```

Asynchronous DECnet Problems

2. Perform the following step for dynamic asynchronous DECnet only. Otherwise, go to step 3.

Use the following commands to install DYN SWITCH and VIRTUAL terminals:

- a. Do the following on both nodes:

```
$ INSTALL:=$SYS$SYSTEM:INSTALL
$ INSTALL/COMMAND
INSTALL> CREATE SYS$LIBRARY:DYN SWITCH/SHARE-
_INSTALL> /PROTECT/HEADER/OPEN
INSTALL> EXIT
$
```

- b. To install virtual terminals do the following on both nodes:

```
$ MCR SYSGEN
SYSGEN> CONNECT VTA0/NOADAPTER/DRIVER=TTDRIVER
SYSGEN> EXIT
$
```

3. Set the terminal port characteristics correctly.

Note

This example assumes a dialup connection.

- a. For nondialup connections, remove the MODEM and NOHANGUP parameters.
- b. For static asynchronous DECnet, execute the following command to insure the terminal characteristics are set up properly.

```
$ SET TERMINAL/PERMANENT/PROTOCOL=DDCMP/NOTYPE_AHEAD/MODEM-
_$ /NOHANGUP/EIGHT_BIT/NOAUTOBAUD/SPEED=xxx TTA0:
```

- c. For dynamic asynchronous DECnet, execute the following command to insure the terminal characteristics are set up properly.

```
$ SET TERMINAL/PERMANENT/NOTYPE_AHEAD/MODEM/DISCONNECT-
_$ /EIGHT_BIT/NOAUTOBAUD/SPEED=xxx tta0:
```

B. Check the DECnet software to make sure that the circuit is set up properly in NCP.

1. For *static asynchronous DECnet* only, run NCP and display the characteristics with the following command:

```
NCP> LIST LINE tt-0-0 CHARACTERISTICS
NCP> LIST CIRCUIT tt-0-0 CHARACTERISTICS
```

2. Use the following command to set the line state, receive buffers (4 is the recommended default), and line speed:

```
NCP> DEFINE LINE tt-0-0 STATE on RECEIVE BUFFERS 4-
_NCP> LINE SPEED 2400
```

3. Set the circuit state on using the following command:

```
NCP> DEFINE CIRCUIT tt-0-0 STATE ON
```

C. Make sure the communications link is functioning properly.

1. For dialup problems, see "Dialup Problems" in this chapter.

Note

Modems that perform automatic error correcting often do not work with asynchronous DDCMP lines. If a line works without DECnet but fails when you start DECnet, check to see if the modem error correcting is enabled, and turn this option off if possible.

2. For systems that are hardwired, or directly cabled, or both, set the circuit state on, and perform loopback tests (see NCP Loopback Tests) as follows:

- For static asynchronous DECnet execute the following commands:

```
NCP> SET CIRCUIT tt-0-0 all
NCP> LOOP CIRCUIT tt-0-0
```

The loopback tests shows if a bad cable, modem, circuit, or device exists.

- For dynamic asynchronous DECnet, first, use the following command to temporarily make the line a static asynchronous DECnet line for the loopback testing:

```
$ SET TERMINAL/PERMANENT/PROTOCOL=DDCMP/NOTYPE_AHEAD-
_$ /MODEM/NOHANGUP/EIGHT_BIT/NOAUTOBAUD/SPEED=xxxx TTA0:
```

Then, run NCP and use the following commands to set the line and circuit characteristics and perform loopback testing:

```
NCP> SET LINE tt-0-0 STATE on RECEIVE BUFFERS 4-
_NCP> LINE SPEED 2400
NCP> SET CIRCUIT tt-0-0 STATE on
NCP> LOOP CIRCUIT tt-0-0
```

For more information, also see "Circuit State Problems" in this chapter.

Babbling Device

Babbling Device

Symptoms

Users perceive very slow response time for network operations.

Explanation

This symptom indicates a LAN problem. A device on the network has a hardware problem, causing it to send out large amounts of data to the local area network. Such a device is called a *babbler* or *babbling device*.

Troubleshooting Strategy

To solve this problem, determine the LAN segment on which the babbling device exists, and the physical location of the babbling device on that segment. After you determine the location of the device, use hardware diagnostic techniques to determine the cause of the problem.

Troubleshooting Procedure

1. Determine the physical location of the babbling device using LTM or DECelms. If LTM or DECelms are not available at your site, go to step 3.
 - If your site uses LTM, use the menus to determine the physical address of the babbling device.

Under the heading *Top Ten Talkers*, LTM displays the addresses of the devices it recognizes that are generating the most traffic. After you have the physical address of the babbling device, you can use DECelms to determine the physical location of the babbling device.

Note

LTM displays only the addresses of devices it recognizes. If the babbling device is generating corrupt information (such as a long stream of preamble) for its own name and address, you may have trouble determining the physical address or location. In this case, you can use LTM to capture bad packets. Using the information on bad packets, you may be able to deduce the address of the problem device.

If you cannot determine the address of the babbling device using LTM, go to step 3.

-
- If your site uses DECelms, and you know the physical address of the babbling device, do the following to locate the physical location of the babbling device:
 - a. Run DECelms and examine the node's address in the bridge's forwarding database using the following commands:

```
ELMS> USE bridge-id  
ELMS> SHOW ADDRESS node-address
```

DECelms displays the last line to recognize communication from the physical address specified.

- b. To further isolate the location of the babbling device, use the following command on the next bridge to display the line on that bridge that last saw the physical address.

```
ELMS> USE bridge-id  
ELMS> SHOW ADDRESS node-address
```

- c. Continue querying the bridges in this way until you isolate the babbling device to a specific LAN on the network.
2. After you determine the LAN segment on which the problem device is located, do the following:
 - a. Ask users on that segment if there have been any new devices added to the LAN, or if they have been having any problems with devices on the LAN.
 - b. Check these devices for correct operation.
 - c. Ask the system manager of the system causing the problem to disconnect the faulty device.
 3. If you do not have LTM or DECelms at your site, or if you cannot determine the physical address of the babbling device, trace the problem to the babbling device by using the following steps:
 - a. Collect information on the amount of data transferred from and received by each node in the network for a given time.
 - b. Compare the data for all the systems.
 - c. Use an Ethernet scope or LAN analyzer to determine the source of the problem, or disconnect each individual node until you find the babbling device.

Start disconnecting nodes that are losing connections and having significant performance problems across all protocols.

Broadcast Storm

Broadcast Storm

Symptoms

Only the host computers having the highest performance CPUs and Ethernet controllers can gain access to the network. Host response time becomes slow and utilization percentage becomes very high. In the DECelms Line Counters display, the Transmit Multiple Collisions and Collision Limit Exceeded values become very high. The LAN utilization percentage values displayed by the LAN Bridge 200 Line Monitor and LAN Traffic Monitor (LTM) rise toward 100%.

Explanation

This symptom indicates a LAN problem. Hosts may be using protocols that overuse the Ethernet broadcast address. By default, the bridges are required to pass these messages to every segment in the network, even to segments where there is no possible recipient. Every host in the network must read all the frames sent to the broadcast address.

Causes of broadcast storms include:

- Protocol problems — Certain protocols rely on the use of the Ethernet broadcast address instead of using multicast addresses, for example, ARP broadcasts.
- Configuration problems — Many different versions of both 4.2BSD and 4.3BSD UNIX, and many vendor's operating systems based on these different versions exist. Some operating systems perform differently from other operating systems when they receive broadcast packets they do not recognize as broadcasts. The different ways operating systems handle broadcast packets can cause problems.

For example, when hosts using 4.3BSD UNIX are placed in a network of 4.2BSD UNIX hosts, a major "ARP storm" often occurs. The 4.2BSD hosts have gateway software enabled by default and consider any IP broadcast address other than zeros as an unknown address. When a 4.3BSD host using the new broadcast address of ones is added to the network, the 4.2BSD hosts simultaneously broadcast an ARP lookup request whenever the new 4.3BSD host uses the IP broadcast address of ones. The resulting broadcast storm can have a crippling effect on network performance.

Troubleshooting Strategy

Follow this strategy to troubleshoot a broadcast storm:

1. Determine the protocol that is causing the storm.
2. Contain the problem by adding filters for the protocol in all the LAN Bridge 200 models.
3. Identify the segment where the storm originated and resolve the problem.

Troubleshooting Procedure

To troubleshoot a broadcast storm, follow these steps:

1. Use LAN Traffic Monitor (LTM) to determine the protocol that is causing the broadcast storm.
 - a. From the LTM Main Menu, select #2, Node, Type and Multicast Traffic Displays.
 - b. From the Node, Type and Multicast Traffic Displays Menu, select #4, Multicast Traffic by Type Display. LTM prompts you for a multicast address.
 - c. Enter the broadcast address, FF-FF-FF-FF-FF-FF. LTM displays the following:

```
MC Address : FF-FF-FF-FF-FF-FF  Broadcast      Type Count : 3
Type Field 06-00 Xrx_NSIDP  Frame Count      2820          7.51%
Type Field 08-00 DOD_TCPIP  Frame Count     10272         14.09%
Type Field 08-06 TCPIP_ARP  Frame Count     10272         67.41%
```

In this case, the broadcast storm is caused by TCP/IP, which has the Protocol Type 08-00, and by the Address Resolution Protocol (ARP), which is used by TCP/IP for locating addresses. ARP has the Protocol Type value 08-06. According to the display, ARP is generating 67.41% of the broadcast traffic on the network.

2. Contain the problem by adding filters for the Protocol Types 08-00 and 08-06 to the protocol databases of the LAN Bridge 200 models using DECelms. These filters instruct the bridges to filter (discard) frames that contain TCP/IP or ARP protocol information. The effect is to isolate the broadcast storm as much as possible, perhaps even to a single segment, if a LAN Bridge 200 connects the segment where the broadcast storm originated.

- a. If the LAN Bridge 200 models do not have a password set, enter the following commands:

```
ELMS> USE KNOWN BRIDGES
ELMS> ADD PROTOCOL 08-00 DISPOSITION FILTER
ELMS> ADD PROTOCOL 08-06 DISPOSITION FILTER
```

- b. If the LAN Bridge 200 models have the same password, enter the following commands, where *password* is the password:

```
ELMS> USE KNOWN BRIDGES
ELMS> ADD PROTOCOL 08-00 DISPOSITION FILTER PASSWORD password
ELMS> ADD PROTOCOL 08-06 DISPOSITION FILTER PASSWORD password
```

- c. If the LAN Bridge 200 models have different passwords, repeat the following commands for each bridge, where *bridge-name* is the name of the bridge and *password* is its password:

```
ELMS> USE bridge-name
ELMS> ADD PROTOCOL 08-00 DISPOSITION FILTER PASSWORD password
ELMS> ADD PROTOCOL 08-06 DISPOSITION FILTER PASSWORD password
```

Broadcast Storm

- d. Because LAN Bridge 100 and LAN Bridge 150 models do not support protocol filtering, you must instruct these bridges to filter frames sent to the broadcast address by entering the following commands, where *bridge-id* is the name of the target bridge. If the target bridge is a LAN Bridge 150 that has a password set, you must include **PASSWORD** and the bridge's password.

```
ELMS> USE bridge-id
ELMS> ADD ADDRESS FF-FF-FF-FF-FF-FF DISPOSITION FILTER-
_ELMS> PASSWORD password
```

3. Locate the segment where the storm originated and resolve the problem.

- a. From the LTM Main Menu, select #2, Node, Type and Multicast Traffic Displays.
- b. From the Node, Type, and Multicast Traffic Displays Menu, select #6, List of Nodes Using Protocol Type. LTM prompts you for a Protocol Type value.
- c. Enter 08-00 or 08-06 to display a list of the hosts using TCP/IP or ARP. LTM displays the addresses of the hosts using the protocol.
- d. Follow the procedure described in the Babbling Device example to find the segment where the broadcast storm originated.
- e. Resolve the problem, or at least ensure that it is isolated to a single segment. If the bridge connecting the segment is a LAN Bridge 200, enter the following commands to ensure that it has the appropriate protocol entries in its protocol database. If it does not, add them as described in step 2c.

```
ELMS> USE bridge-id
ELMS> SHOW PROTOCOL 08-00
ELMS> SHOW PROTOCOL 08-06
```

If the bridge connecting the problem segment is a LAN Bridge 100 or LAN Bridge 150 model, ensure that the bridge has an entry for the broadcast address in its forwarding database:

```
ELMS> USE bridge-id
ELMS> SHOW MULTICAST ADDRESS FF-FF-FF-FF-FF-FF
```

4. Finally, remove any unnecessary protocol or address entries that you added when you were first trying to contain the problem. The following commands remove a protocol entry from the protocol database of a LAN Bridge 200:

```
ELMS> USE bridge-name
ELMS> REMOVE PROTOCOL 08-00 PASSWORD password
```

The following commands remove an address entry from the forwarding database of a bridge. (You do not need to include the password if the target bridge is a LAN Bridge 100 model or if it does not have a password set.)

```
ELMS> USE bridge-id
ELMS> REMOVE ADDRESS FF-FF-FF-FF-FF-FF PASSWORD password
```

Recommendations

- You can minimize broadcast storms by adding protocol filters to isolate protocols that misuse the broadcast address. Use LTM and DECelms to study the protocols used on each segment in your extended LAN and use filters to contain protocols to the LAN segment where they are used. (You can do the same for multicast addresses also.)

For example, if one segment in an extended LAN supports a classroom of nodes using AppleTalk, you can add a protocol entry for AppleTalk with the disposition FILTER in the bridge connecting the segment. The filter prevents the AppleTalk broadcasts from entering the extended LAN, where there are no nodes using AppleTalk. This prevents the broadcasts from taking up valuable network bandwidth and host processing power.

- Because of the potential problems caused by the different ways operating systems handle broadcast packets, make sure that you install recent releases of the respective operating systems on your network. If you cannot upgrade a system, see an expert about possibly creating a patch to prevent potential broadcast storms.

Circuit State Problems

Circuit State Problems

Symptoms

The system displays one of the following messages when you use the NCP command, `SHOW CIRCUIT circuit-id`:

```
NCP> SHOW CIRCUIT circuit-id
Circuit on-starting
Circuit on-synchronizing
Circuit off-synchronizing
```

Explanation

This symptom indicates a DECnet-VAX node problem that sometimes manifests itself as a LAN or WAN problem initially.

- A. Circuit on-starting is generally a normal condition indicating that the circuit is ready to begin a node initialization sequence. Circuit on-starting is only a problem if an adjacent node is connected, or should be connected to the circuit that is on-starting.

When it is not a normal condition, circuit on-starting indicates a problem with point-to-point (non-Ethernet) links, including any of the following:

- The remote node is not running.
 - The line is not connected, or there is a bad connection on the cables for the devices.
 - The modem is not running.
 - The circuit on the remote node is in the off state.
- B. Circuit on-synchronizing means the node initialization sequence between two adjacencies is failing, and usually indicates a routing-related problem. Generally, circuit on-synchronizing can be caused by any of the following:
- DECnet event 4.2, Node out-of-range packet loss
 - DECnet event 4.3, Oversized packet loss
 - DECnet event 4.6, Verification reject
 - The circuit is on and the executor state is off
 - The circuit is on and the line is off
- C. Circuit off-synchronizing indicates a hardware problem.

Troubleshooting Strategy

- A. For circuit on-starting, do the following:
1. Check the configuration.
 2. Check that the remote node is running, and its lines and circuits are on.
 3. Check that the local node's lines are on.
 4. Use loopback tests.

- B. For circuit on-synchronizing, do the following:
 1. Enable event logging.
 2. Check for class 4 DECnet events, and correct as described in step B.
 3. Make sure that the executor state is on.
 4. Make sure that the lines and circuits are on.
- C. For circuit off-synchronizing, check the device hardware manuals to resolve the hardware problem.

Troubleshooting Procedure

A. For circuit on-starting problems, do the following:

1. Check the configuration that should exist.
2. Check with the system manager of the remote node to see if the system is running, and if the circuits and lines are turned on.
3. Run NCP and use the following command to make sure that the local lines are turned on:

```
NCP> SHOW KNOWN LINES
```

4. If the local lines are not on, use the following NCP command to turn them on:

```
NCP> SET LINE line-id STATE ON
```

5. Use loopback tests to determine where the physical problem exists. The problem could be at any one of many locations. Systematically work across the physical connection, looping back to test at successively more remote points, using the loopback tests described in Chapter 4.
 - For controllers, use controller loopback tests.
 - For distribution panels or cables, use a loopback connector.
 - For modems, use modem local loopback tests.
 - For communications lines, use modem remote loopback tests.

B. For circuit on-synchronizing problems, do the following:

1. Make sure that OPCOM is running, and use the following DCL command to enable event logging:

```
$ REPLY/ENABLE=NETWORK
```

2. Check for class 4 DECnet events.
 - DECnet event 4.2, Node out-of-range packet loss, means that the remote node's address is greater than the local executor's MAXIMUM ADDRESS parameter. See "Node Out of Range Packet Loss" problem in this chapter, for more information on this problem.
 - DECnet event 4.3, Oversized packet loss, means that the routing layer discarded a packet because the packet was too large to forward to an adjacent node. The solution is to ensure that all nodes in the network have the same executor buffer size, then to stop and restart the network.

Circuit State Problems

```
NCP> DEFINE EXECUTOR BUFFER SIZE value
NCP> SET EXECUTOR STATE OFF
NCP> EXIT
$ @STARTNET.COM
```

- DECnet event 4.6, Verification reject, means that a problem exists with the transmit or receive passwords. The solution is to make sure that the transmit and receive passwords match. See "Verification Reject" in this chapter for more information on this problem.
3. If the circuit is on and the executor state is off, use the following NCP command to turn the executor state on:

```
NCP> SET EXECUTOR STATE ON
```

4. Use the following command to restart DECnet, if necessary:

```
$ @STARTNET.COM
```

5. If the circuit is on and the line is off, use the following NCP command to turn the line on:

```
NCP> SET LINE line-id STATE ON
```

- C. **If the circuit is off-synchronizing, the device has a hardware problem.** Follow the procedures in the appropriate hardware manuals or refer the problem to Customer Services.

Connect Failed, Access Control Rejected

Symptoms

ULTRIX system users receive an error message such as the following when attempting any network operation:

```
connect failed, access control rejected
```

Explanation

This symptom indicates a DECnet–ULTRIX host problem involving the session layer. It can be caused by any of the following:

- Incorrect user-supplied access control information
- Incorrect proxy access set up
- Invalid account specified for the object
- Incorrect or nonexistent proxy database or proxy accounts
- Restricted proxy access on either the local or the remote host

Note

This message may not indicate a problem; the host may be restricting incoming network access for security reasons.

To troubleshoot this problem effectively, you need to understand the order of access control for DECnet–ULTRIX.

ULTRIX systems check first for explicit, user-supplied access control information. If user-supplied information does not exist, the ULTRIX system checks for proxy access information. If proxy access information does not exist, the ULTRIX system checks for default access information.

Figure 5–1 shows the order of access control for ULTRIX systems in more detail.

Connect Failed, Access Control Rejected

Figure 5-1 ULTRIX Access Control

Note

NCP parameters determine the type of proxy access permitted, if any, for objects on a host or for the host itself. NCP information specified for an object always supersedes the executor information. This is true for proxy as well as default account information.

Connect Failed, Access Control Rejected

Proxy access for the executor is determined according to the parameters shown in Table 5–2.

Table 5–2 NCP Executor Proxy Access Parameters

Executor Parameter	Function
INCOMING PROXY DISABLED	Ignores all incoming proxy requests, and instead, relies exclusively on access control information supplied in the connect requests to validate the logical link.
INCOMING PROXY ENABLED	Invokes the appropriate proxy, based on the source user, source node, and supplied access control information, if any. This is the default.
OUTGOING PROXY DISABLED	Specifies that proxy login is not requested on any outgoing logical links.
OUTGOING PROXY ENABLED	Specifies that proxy login is requested on outgoing logical links. This is the default.

Troubleshooting Strategy

Before you attempt to solve this problem, do the following:

- Determine the type of access control the user specified when trying to access the remote host or node.
 - Determine the objects the user tried to access.
- A. For user-specified access control problems, do the following:
1. Check the user-specified access control information.
 2. Make sure that the account the user wants to access exists on the remote host, and create one if appropriate.
 3. Modify the password information for the remote account.
- B. For proxy access control problems, do the following on the source (local) and target (remote) hosts:
1. On the source (local) host, make sure that PROXY OUTGOING is enabled.
 2. On the target (remote) host, do the following:
 - a. Make sure that PROXY INCOMING is enabled
 - b. Make sure that the `/etc/dnet_proxy` file has the correct, case-sensitive entries.
 - c. Make sure that the local host is defined in the remote host's DECnet database.
- C. For default access control problems, do the following:
1. If the requested object on the remote host has an account associated with it, check the `/etc/passwd` file to make sure that the account exists.
 2. If the requested object on the remote host does not have an account associated with it, create an account for the object and modify the object to ensure that the object has an account associated with it.

Connect Failed, Access Control Rejected

Troubleshooting Procedure

Before you try to solve this problem, do the following:

Use the following list to help determine which step to use to solve the problem:

- If the user specified explicit access control information when trying to access the remote host, use step A to solve the problem.
- If the user tried to access the remote host using proxy access, use step B to solve the problem.
- If the user did not try to access the remote host using proxy access, use step C.
- If you do not know if the user tried to access the remote host using proxy access, use step B first, and continue with step C if necessary to solve the problem.

Note

To help solve this problem faster, you can try to connect to a different object on the remote host. If the connection succeeds when directed toward the new object, then the problem is probably object-specific and not related to proxy access. You can focus your efforts on ensuring that the object on the remote host is set up properly.

A. For user-supplied, explicit access control problems, do the following:

1. Try to log in to the remote host with the same access information the user tried. If you cannot log in to the host, the access information is probably incorrect.
2. Log in to an account on the remote host.
3. Log in to the superuser account.
4. Look at the password file, using the following command:

```
# cat /etc/passwd
```

5. If the `/etc/passwd` file does not list the account the user tried to access, use the following command to add the account if a new account is appropriate:

```
# adduser
```

The `adduser` command prompts you for the following information about the new account for the new user:

- Login name
- Full name
- Login group (default is `[users]`)
- Other groups
- Parent directory (default is `[/usr/users]`)
- Password

Connect Failed, Access Control Rejected

The `adduser` command adds the new user account to the `/etc/passwd` file, and sets up a home directory for the new user containing the files `.cshrc`, `.login`, and `.profile`.

Note

In ULTRIX versions prior to 4.0, the `adduser` command does not prompt for a password for the user, so be sure to use the `passwd` command to specify a password for the user if you want to prevent unauthorized access to the account.

6. If you are using an ULTRIX version prior to 4.0, use the following command to specify a password for the account, substituting the new user's password for `newpassword`:

```
# passwd username
New password: newpassword
Retype new password: newpassword
```

The characters you type for the new password are not displayed on the screen.

7. Try to log in to the user's account with the new password.
8. **For proxy access control problems, do the following:**
 1. On the source (local) host do the following:
 - a. Log in to an account on the local host.
 - b. Log in to the superuser account.
 - c. Run `ncp`, and use the following command to display the executor characteristics and determine if proxy outgoing is enabled:

```
ncp> show executor characteristics
```
 - d. If proxy outgoing is not enabled, use the following `ncp` commands to enable it in both the volatile and permanent databases:

```
ncp> set executor proxy outgoing enabled
ncp> define executor proxy outgoing enabled
```
 2. On the target (remote) host, do the following:
 - a. Run `ncp`, and use the following commands to display the executor characteristics and determine if proxy incoming is enabled:

```
ncp> show executor characteristics
```
 - b. If proxy incoming is not enabled, use the following `ncp` commands to enable it in the both volatile and permanent databases:

```
ncp> set executor proxy incoming enabled
ncp> define executor proxy incoming enabled
```
 - c. Display or edit the `/etc/dnet_proxy` file to make sure that it has the correct, case-sensitive entries for the local host and user.
 - d. If entries are missing from the `/etc/dnet_proxy` file or are incorrect, edit the `/etc/dnet_proxy` file, using the following format for the entries:

Connect Failed, Access Control Rejected

```
source::user local_user
```

- e. Run `ncp` and use the following commands to make sure that the source (local) host is defined in the volatile and permanent DECnet databases:

```
ncp> set node hostname address aa.nnn  
ncp> define node hostname address aa.nnn
```

9. Do the following to resolve problems due to default access control information:

- a. Log in to an account on the remote host.
- b. Run `ncp`, and use the following command to display the default user account:

```
ncp> show object object_name characteristics
```

- c. If the object does not have an account associated with it, use the following command to create an account if a new account is appropriate:

```
# adduser
```

- d. Run `ncp`, and modify the object using the following command:

```
ncp> set object object_name default user account_name
```

Connect Failed, Unrecognized Object

Symptoms

Users on an ULTRIX system receive the following message when trying to access a remote host:

```
connect failed, unrecognized object
```

Explanation

This is a DECnet-ULTRIX host problem involving the session layer on the remote host. The connect failed message occurs when the object requested is not defined in ncp, or if the requested object has file protection problems.

Troubleshooting Strategy

1. Check to see if the object is defined in ncp.
2. If the object is not defined, define it.
3. Check to see if the file specified for the object exists.
4. If the file for the requested object does not exist, create it.
5. Make sure that the protection specified for the file is correct.

Troubleshooting Procedure

1. Run ncp and use the following command to see if the object is defined on the remote host.

```
ncp> tell remote-node-id show known objects
```
2. If the object is not defined, log in to the superuser account and run ncp to define the object using the following command and additional parameters as required:

```
ncp> set object object-id
```
3. If the object is defined, run ncp and use the following command on the remote host to see if the object has a file specified:

```
ncp> tell remote-node-id show object object-id characteristics
```
4. Use the following command to see if the file specified for the object exists.

```
# ls -l
```
5. If the file for the requested object does not exist, create the file.
6. Ensure that the protection on the specified file is correct. Generally, world execute access is required for most objects.
Use the following commands to set the file protection:

```
# chmod a+x /usr/etc/fal
```

Connect Failed, Unrecognized Object

Use the following commands to ensure that the directories above the file (including the root (/) directory) also have the correct file protection to allow access:

```
# cd /usr  
# ls -ld etc  
# chmod a+x etc
```

Connection Timed Out

Symptoms

TCP/IP network users receive the following error when attempting any TCP/IP-based network operation:

```
connection timed out
```

Explanation

This message occurs when the TCP software attempts to make a connection to the destination host, and does not receive any packets in response to the packets it sends. Most often, connection timed out is the result of a problem on the source or destination hosts, not a problem on a host on the path between the two.

Potential causes for this problem are as follows:

- The destination host is not running.
- A host's broadcast address or address mask is incorrectly defined in its `/etc/rc.local` file.
- The local host does not have its IP address properly defined in the `/etc/hosts` file.
- ARP entries on the source or destination host are inaccurate.
- A cabling problem exists.
- A LAN problem exists.
- An intermediate router is not running, but the routing protocols have not discovered this yet.
- A WAN problem exists.

Troubleshooting Strategy

- A. Determine whether the problem is on the source host, the destination host, or a host on the path between the two.
- B. On the problem host, check for problems with the following:
 - Broadcast address
 - `/etc/hosts` file
 - Hardware
 - ARP entries
 - LAN connections
 - WAN connections

Connection Timed Out

Troubleshooting Procedure

A. Determine whether the problem is on the source host, the destination host, or a host on the path between the two.

1. Use the ping command to determine if the destination host is running:

```
% ping hostname
```

- If the ping command returns the message, "hostname is alive," the destination host is operational. The destination host may have just been coming up when the user tried to reach it before, or the problem may be transient. Try the original network operation again.
- If the ping command returns the message, "no response from hostname," continue with the next step.

2. Use the ping command to determine if the source host can reach other hosts on its subnet.

```
% ping hostname
```

- If the source host can reach other hosts on its subnet, go to the destination host, if possible. Use the ping command to see if the destination host can reach other hosts on its subnet.

If you cannot physically check the destination host, or call someone on the destination host to check its availability, then work with your local network administrator to try to determine the status of the destination host.

- If the destination host can reach other hosts on its subnet, then the problem may involve an IP router between the source and destination hosts. Trace the routing path using netstat (see netstat Routing Path Trace Procedure) or traceroute (see Chapter 3) to locate the problem host, and go to step B.
- If the destination host cannot reach other hosts on its subnet, then the problem is on the destination host. Go to the destination host and continue with step B.
- If the source host cannot reach other hosts on its subnet, then source host is the problem host. Go to step B.

B. Perform the following steps for the problem host.

The problem host may be the source host, the destination host, or any host on the path between the source and destination hosts, as you determined in step A or through tracing the path.

The following steps use the term "local" to refer to the host on which you perform the action required. The term "remote" refers to any host you try to reach using the action.

1. Confirm that the broadcast address and address mask for the local host are properly setup in the /etc/rc.local file, and that the network device is properly configured.

If you are not sure what the broadcast address and address mask is for the local host, check with the local network administrator, and make any changes necessary in the /etc/rc.local file.

- a. Use the following command to display the configured network devices:

```
# netstat -i
```

- b. If the network device is not configured, configure it using the `/etc/ifconfig` command, using the following example as a guideline:

```
# /etc/ifconfig qe0 `bin/hostname` broadcast 16.0.255.255  
netmask 255.255.0.0
```

2. Make sure the local host's `/etc/hosts` file has the correct IP address for the local host.

If the IP address is incorrect, the local host can reach other hosts, but other hosts that try to reach the local host receive connection timed out.

3. Make sure the cabling from the local host to the network is intact and properly connected.
4. Use the following `netstat` command to determine whether any input or output errors exist.

```
% netstat -i
```

Input errors indicate that a host or hosts are sending bad packets. Most likely, the problem is a hardware error on the host sending the bad packets. Use a protocol analyzer or LTM to determine which host is sending the bad packets.

Output errors indicate a hardware problem on the problem host. Use the following `uerf` command to display errors, then call Customer Services:

```
# uerf -R
```

Note

If the remote host is connected to the local host through a LAN connection, perform steps 5 through 8.

If the remote host is connected through a WAN connection, go to step 9.

5. If the remote host is on the problem host's LAN, use the following `arp` command to delete the entry for the remote host from the translation tables:

```
# arp -d hostname
```

6. Use the `ping` command to try to reach the remote host, as follows:

```
# ping hostname
```

Because the translation tables no longer contain an entry for the remote host, the `ping` command generates an ARP request for the remote host to reply with its Ethernet address.

- If the remote host is available, it responds with its Ethernet address and the message, "hostname is alive," indicating it is reachable through IP. Try the original network operation again.

Note

If you are performed this step on a host on the routing path, continue tracing the routing path. If you encounter another problem host, go to that host, and repeat step B for that host.

Connection Timed Out

- If the remote host is not available, the ping command returns the message, "no response from hostname." Continue with the next step.

Note

The recommendations section for this problem provides additional information on ARP-related problems.

7. Verify that the local host's software connection to the network is working properly by using the ping or rlogin command to see if the local host can reach other hosts on the local network.
8. If you can not get to other directly-connected hosts, a LAN problem such as LAN segmentation, a babbling device, or a broadcast storm may exist. Use tools such as LAN Traffic Monitor, NMCC/VAX ETHERnim, or DECMCC Management Station for ULTRIX to isolate the problem, and see the troubleshooting procedures in this chapter for "LAN Segment Communication Problem," "Babbling Device," and "Broadcast Storm."

Note

If the remote host is connected to the local host through a WAN connection, perform steps 9 through 12.

9. If the furthest host you were able to reach when tracing the routing path is connected through a point-to-point link (WAN), use the following netstat command to display errors on that host:

```
% netstat -i
```

- If the netstat command displays no errors, the network connections are working properly, but the remote host may be down.
 - If the netstat command displays input or output errors, a host modem, wire, or cable is sending bad packets or corrupting packets. Continue with the next step, performing modem loopback tests to determine if the problem is one of the following:
 - Local or remote hardware
 - Common carrier circuit between hosts
 - Cabling between the modem and the interface
 - Local or remote modem
10. Perform local loop tests on the modems at both the local and remote ends. If either the local or remote modems fail, this is the source of the connection timed out problem. Replace the failing modem, and try the original network operation again. Otherwise, continue with the next step.
 11. Put the modem on one end in remote loop mode and do a remote loop self test to test the modem's operation with the common carrier circuit. If this test succeeds, go to step 12.

If the local loop self test had succeeded both the local and remote ends, and the remote loop self test fails, the common carrier circuit is out of order.

Call the common carrier for repair service.

12. Use a breakout box to make sure that the cables are connected and the signaling is correct on both the local and remote ends, and that both ends can transmit and receive data. Repair any broken cables.

Recommendations

- In solving this problem, you may find that an intermittent or transient problem may have caused the original connection timed out message. If your site uses tools that record historical data, check to see if any thresholds were reached or surpassed. These threshold values might have caused the connection timed out message.
- ARP tables can contain inaccurate entries, due to the way some systems perform ARP cache timeouts. Inaccurate ARP entries can occur in the following cases:
 - When a host's Ethernet interface is replaced
 - When DECnet starts on a host
 - When a system running DECnet reboots, but does not restart DECnet

To help isolate inaccurate ARP entries, check to see if various hosts on the same Ethernet can reach another host on the Ethernet. For example, if Host X can reach Host Y, but Host Z cannot reach Host Y, check the ARP entries for Host Z.

To solve problems due to inaccurate ARP entries, remove the old ARP entry from the translation tables using the `arp -d` command.

Note

Problems can also occur with ARP entries on hosts running both DECnet and TCP/IP, if the software is not started in the proper order. If a host is running both DECnet and TCP/IP, make sure the DECnet software starts first, so that IP never propagates the non-DECnet Ethernet address through ARP.

When Phase IV DECnet starts, it modifies the Ethernet hardware address with a six-octet DECnet address, even if another protocol is already started and is using the address. DECnet must modify the Ethernet address to be able to function on the Ethernet controller.

If a host is running only TCP/IP, its ARP entries are based on the unaltered Ethernet address. If the host then starts DECnet (which changes the Ethernet address), all the existing ARP entries become incorrect.

DECnet Ethernet addresses start with AA-00-04-00. The last two octets of the Ethernet address are the DECnet node address.

Device Not Mounted

Device Not Mounted

Symptoms

Users receive the following error message when attempting to perform any network operation:

```
Device not mounted.
```

Explanation

This is a DECnet-VAX problem is related to the local node. When DECnet starts, SYSGEN loads the necessary drivers, and creates and mounts the NET0 device. However, this message shows that DECnet is not running, because a network application attempted to open the NET0 device to perform a DECnet operation, but the device is not mounted.

Troubleshooting Strategy

To resolve this problem, start DECnet. If starting DECnet does not resolve the problem, check to see if NETACP is running.

Troubleshooting Procedure

1. Use the following command to determine if the NET devices are loaded:

```
$ SHOW SYSTEM
```

If NETACP is one of the process names listed, the devices are loaded and DECnet is running.

2. If the NET devices are not loaded, use the following command to start DECnet:

```
$ @SYS$MANAGER:STARTNET
```

3. If the previous command returns an error, it may be due to problems with the LOADNET.COM file. Do the following to resolve LOADNET.COM problems:

- a. Make sure that the LOADNET.COM file exists.

The STARTNET.COM file calls the LOADNET.COM file, and, if the LOADNET.COM file does not exist, the STARTNET procedure fails.

- b. Run NCP, and use the following command to stop the network:

```
NCP> SET EXECUTOR STATE OFF
```

- c. Exit NCP, and use the following command to display command lines and data lines from the STARTNET procedure:

```
$ SET VERIFY
```

Resolve any problems indicated by the command and data lines from the STARTNET procedure.

- d. Use the following command to restart the network:

```
$ @SYS$MANAGER:STARTNET
```

4. If starting DECnet does not resolve the problem, something more complex is occurring. Check to see if NETACP is running, using the following command:

```
$ SHOW SYSTEM
```

Look for the NETACP process. NETACP must be running for the network to be running. If NETACP is not running, check to see if the image has been corrupted, or whether you have defined a logical name for NET. If you have a logical name definition for NET, the network cannot start because DECnet uses NET as a device name.

Other factors that may interfere with NETACP include the following:

- Insufficient quotas
- Incorrect system parameters, such as the number of process slots

Dialup Problems

Dialup Problems

Symptoms

Users cannot dial up to a remote node.

Explanation

This symptom indicates a cross-category problem.

The failure of dialup connections may be due to a problem with any of the following:

- Local end
- Remote end
- Telephone lines
- Modems
- Connecting cables

Troubleshooting Strategy

To solve this problem, evaluate and repair each of the potential problem areas in this order: first the local end, then the remote end, and finally the telephone lines.

Make sure that the set up parameters (such as speed, parity, modem control, and so forth) on the local and remote ends are properly defined. Ensure that the telephone lines are operational.

Troubleshooting Procedure

1. On the local end, make sure that the speed, parity, bits, modem control, flow control, and other terminal characteristics are set up properly for the type of modem you have.

Note

The following example shows how to set these parameters for a VAX/VMS system using a DF242 modem.

```
$ SET TERMINAL/PERMANENT/MODEM/DIALUP/HANGUP/SPEED=xxxx tta0:  
$ SET HOST/DTE tta0:  
REM-I-TOEXIT, connection established, type ^\ to exit  
  ^B  
Ready
```

Note

If the process fails, check the settings and cabling from the modem to the device. If a response other than "Ready" comes back from the modem, the problem is probably with the modem. Consult the modem manual for further information on how to resolve this problem.

Dialup Problems

2. Dial the number to the remote node.

If successful, the following message is displayed:

Attached (Speed:2400)

- If you get the "Attached" message, but not a login prompt such as "Username:", check the terminal server or VAX node at the remote end to be sure that the terminal port is set up properly.
 - If you get a message other than "Attached," plug a telephone handset into the local telephone line to check for a dial tone.
 - If you hear a dial tone, the telephone line is working. Continue with step 3.
 - If you do not hear a dial tone, call your local carrier to fix this problem.
 - If you get no message, make sure that the cabling between the local system and the modem is intact, and that the local system and the modem do not have hardware problems. Continue with step 3.
- ### 3. If you get a dial tone, do the following to further isolate the problem:

- a. If the modem has local loopback capabilities, use the local loopback and type characters on the local node or terminal's keyboard.
 - If the characters echo back on the local system, connectivity is intact between the local system and the modem, and the modem parameters are properly set. Go to step 4.
 - If the characters do not echo back, use a loopback connector or breakout box at the back of the local terminal or node, then type characters again.

If characters echo back on the local system now, the local system is operating but the cable between the modem and the DTE is faulty, or there are set up problems (such as bits per character, parity and speed settings) between the DTE and DCE. If the characters do not echo back, the DTE is faulty.
- b. If your modem does not have local loopback capabilities, do the following to isolate the problem:

If either the local or remote end is connected to a terminal server or VAX, temporarily connect a terminal directly to the modem interface so that you have a terminal at each end to use for testing.

Verify that the interface accepts connections by checking to see that the characters you type at one of the terminals are also displayed on the terminal attached to the other end of the modem.

- c. Check the display on the modem.

If the modem is operating properly, the modem displays data terminal ready (DTR) and data carrier detect (DCD) signals. If the modem does not display the DTR and DCD signals, use a breakout box to check the signals between the modem and the interface. The normal progression of RS232 signals between the modem and the interface is as follows.

1. Ring indicator (RI, pin 22) toggles on and off to the DTE.
2. DTE responds with data terminal ready (DTR, pin 20).

Note

When the DTE answers the ring indicator and responds with DTR, RI stops toggling.

3. Modem responds with data carrier detect (DCD, pin 8) and data set ready (DSR, pin 6).
 4. Data passes until the connection ends.
 5. DTE disconnects and the DTR stops.
4. From another handset, dial the local telephone line. If the local telephone rings and you can carry on a conversation, then the telephone line on the local end is good.
If you can not pass voice traffic, or if there is no ring, call your local carrier to fix this problem.
 5. Repeat steps 2 and 3 on the remote node to resolve problems with the remote end.

Host Is Unreachable

Host Is Unreachable

Symptoms

Users on TCP/IP networks receive the following message when trying to access a remote host:

```
host is unreachable
```

Explanation

This is a DECnet–ULTRIX problem, resulting from any of the following:

- The remote host is not available because it is not up and running.
- The local host's routing information for the remote host is incorrect.
- The local host has a problem that prevents it from communicating with any other hosts on the network.
- The destination network or the remote host has a problem that does not prevent the local host from reaching the destination network, but prevents the local host from reaching the destination host on that network.

If the remote host is up and running, but you still cannot reach it, the problem may be caused by any of the following:

- Misconfigured or unconfigured network devices
- Cabling or connection problems
- Improper routing table setup
- Failure to set up routing tables
- Routing daemon problems

Note

This message may not indicate a problem. Routers along the path to the remote host might have security features enabled that prevent you from reaching the remote host.

Troubleshooting Strategy

Assuming that the remote host is up and running, make sure that the following are correct:

- A. Configuration of the network devices on the local host
- B. Routing tables on the local host
- C. Remote host's address-to-name translation on the local host, including the ARP translation for the Ethernet address to IP host name
- D. Configuration of the network devices on the remote host

Troubleshooting Procedure

- A. Make sure that the network devices are configured properly on the local host, using the following steps:

To check the configuration, you need to know the netmask and broadcast address for your network. The `/etc/ifconfig` command sets up the network devices. At system startup, the `/etc/rc.local` file configures the network devices.

1. Use the following command to display the configured network devices:

```
# netstat -i
```

2. If the necessary network device is not configured, configure it using either the `/etc/ifconfig` command or the `netsetup` command as follows:

- To configure the network device with the `/etc/ifconfig` command, use the following example as a guideline:

```
# /etc/ifconfig qe0 `bin/hostname` broadcast 16.0.255.255  
netmask 255.255.0.0
```

This example configures a DEQNA for network 16, with the second octet of the address set for subnet addressing.

- To configure the network device with the `netsetup` command, log in to the superuser account.

For first time configurations, use the following command:

```
#/etc/netsetup install
```

For all existing configurations, use the following command:

```
#/etc/netsetup
```

The `netsetup` program prompts you for information about the remote host, and adds the information you supply to the `/etc/rc.local` file. The changes you make take effect when you reboot the system.

- B. Check the local host's routing tables, remembering that routing can occur through a host-specific route, a route specified for the destination network, or a default route.

See Figure 1–28 for a flow chart illustrating IP routing.

Host Is Unreachable

1. Use the following command on the local host to display the contents of the routing tables.

```
# netstat -r
```

If the routing tables show this routing information	Go to this step
A host-specific or destination network route	Step 2
A default route	Step 3
No route information	Step 4

2. If the routing tables show a host-specific or destination network route for the destination host, use the ping command to see if the IP router specified is reachable.

```
# ping IP_router_name
```

- If you cannot reach the IP router, make sure that the local host's cabling to the network is intact, and do the same for the IP router's cabling to the network.
- If you can reach the IP router, obtain the routing information from the router, and go to the table in step B1. The table in step B1 specifies what to do based on the type of routing information in the routing tables.

3. If the netstat command shows a default route, use the ping command to see if the default IP router is reachable:

```
# ping IP_router_name
```

- If the IP router is not reachable, make sure that the local host's cabling to the network is intact, and do the same for the IP router's cabling to the network.
- If the IP router is reachable, obtain the routing table from the router, and go to the table in step B1. The table in step B1 specifies what to do based on the type of routing information in the routing tables.

4. If no route exists to the destination, add a route.

You can run the routing daemon to add routes automatically or use the route command to add the specific route manually to a router.

Using the routing daemon to add routes automatically

- a. Use the following command to see if the routing daemon (/etc/routed) is running:

```
# ps -aux | grep routed
```

The following example of output from this command shows that `routed` is running in quiet mode. ❶

```
root          77  0.0  0.8  204  88 ? S ❶ 4:42 /etc/routed -q
root          7255 0.0  0.3   40  32 p1 S 0:00 grep routed
```

- b. If the routing daemon is running, but there are still no routes, the local host is not receiving the routing updates.
Check the local host's cabling to the network.
- c. If the routing daemon is not running, but should be, run the routing daemon in quiet mode as follows:

```
# /etc/routed -q
```

- d. Make sure that the `/etc/routed -q` command is in the `/etc/rc.local` file.
- e. Wait a couple of minutes to allow for the routing tables to be filled and try to reach the remote host again.

If you still get the host is unreachable message, go to step B1, and repeat the procedure, now using the updated routing tables.

Using the `route` command to add a route manually

- a. To add a default route to a stable IP router, use the following command:

Note

Use an IP router that is only one hop away.

```
# route add default "ip_router_name" 1
```

- b. Try to reach the remote host again.
 - c. If you still get the host unreachable message, go to step B1, and repeat the procedure.
- C. On the local host, resolve any problems with the ARP entry for the remote host.

1. Clear the ARP table entry for the remote host using the following command:

```
# arp -d hostname
```

2. Use the procedures in the "Unknown Host" problem in this chapter to make sure that the IP address is correct for the hostname.
3. Try to connect to the remote host.
4. If the connection fails, perform the procedure in step A, then go to step 5.
5. If you make changes when you use step A, use the following command to clear the ARP table entry for the remote host:

```
# arp -d hostname
```

Host Is Unreachable

6. Check the ARP table to see if the entry is correct, using the following command:

```
# arp hostname
```

If the arp command shows no translation, the remote host is not responding, and may be down.

- D. Log in to the remote host, and make sure that the network devices on the remote host are properly configured using the procedure in step A.

Insufficient Resources at Remote Node

Symptoms

Users receive the following message when attempting any network operation:

```
%SYSTEM-E-REMRSC, Insufficient system resources at remote node
```

Explanation

This symptom indicates a DECnet-VAX node problem that results from the remote node rejecting a connection because it does not have enough resources to process the request. The message can be caused by the following parameter values:

- SYSGEN parameter, MAXPROCESSCNT
- NCP parameters, MAXIMUM LINKS and ALIAS MAXIMUM LINKS
- AUTHORIZE parameters, MAXJOBS and MAXACCTJOBS

The current settings for these parameters may not be sufficient. For example, the NETACP page file quota may be exhausted, and may need to be modified. The NETACP page file holds the NCP node database. As the number of nodes in the database increases, the page file quota requirements for NETACP increase as well.

Note

This message may not indicate a problem. The parameter values may be set intentionally to disallow network connections beyond a certain number. If someone on the remote node logs off, the local user trying to establish a connection to the remote node may be successful.

Be sure you understand the reason for the current setting before you take any action to solve this problem.

Troubleshooting Strategy

- A. To resolve problems related to the MAXPROCESSCNT parameter, do the following:
 1. Check the number of free process slots.
 2. Check the current value of MAXPROCESSCNT.
 3. Increase the value of MAXPROCESSCNT.
 4. Execute AUTOGEN.COM.
- B. To resolve problems related to the MAXIMUM LINKS and ALIAS MAXIMUM LINKS parameters, do the following:
 1. Check the current values for MAXIMUM LINKS and ALIAS MAXIMUM LINKS on the remote node.
 2. Check the number of links in use at the remote node.

Insufficient Resources at Remote Node

3. Increase the values for MAXIMUM LINKS and ALIAS MAXIMUM LINKS, if necessary.
- C. To resolve problems related to the NETACP page file quota, do the following:
1. Check the current page file quota value.
 2. If the page file quota value is 0, then increase the value, shut down and restart the network.
- D. To resolve problems related to MAXJOBS and MAXACCTJOBS, do the following:
1. Check the current values for MAXJOBS and MAXACCTJOBS specified in the SYSUAF file for the user who received the insufficient resources error.
 2. Use AUTHORIZE to increase the values, if necessary and appropriate.

Troubleshooting Procedure

A. **Do the following to display the current value for MAXPROCESSCNT, and to increase the value, if necessary:**

1. Use the following command on the remote node to determine the total number of process entry slots, as well as the number of free process entry slots:

```
$ SHOW MEMORY
```

The MAXPROCESSCNT value determines the maximum number of process entry slots to be allocated. The default is 32. The maximum is 8192.

The default value for MAXPROCESSCNT normally is sufficient. However, if there have been changes to the system since it was booted, you may need to increase the MAXPROCESSCNT value. For example, if the workload and number of users has changed, you may require a higher number of processes.

2. Edit MODPARAMS.DAT to include the following line, which increases the values for MAXPROCESSCNT:

```
MAXPROCESSCNT=n
```

3. Execute the AUTOGEN.COM file to cause the changes to take effect.

Note

The AUTOGEN.COM command procedure reboots the system. Be sure you really want to reboot the system at this time before you execute the following command:

```
$ @SYS$UPDATE:AUTOGEN GETDATA REBOOT NOFEEDBACK
```

Insufficient Resources at Remote Node

B. Do the following to display the current values for **MAXIMUM LINKS**¹ and **ALIAS MAXIMUM LINKS**², and to increase the values if necessary:

1. Run NCP on the remote node, and use the following command to display the values for **MAXIMUM LINKS** and **ALIAS MAXIMUM LINKS**:

```
NCP> SHOW EXECUTOR CHARACTERISTICS
```

2. Use the following NCP command on the remote node to display the known links:

```
NCP> SHOW KNOWN LINKS
```

3. Count the number of links, and compare that number with the results of the **SHOW EXECUTOR CHARACTERISTICS** command. If the number of known links equals the value for **MAXIMUM LINKS** or **ALIAS MAXIMUM LINKS**, use one of the following commands to increase the maximum links value:

- If the insufficient resources message occurred when the user was connecting to a cluster alias, increase the **ALIAS MAXIMUM LINKS** value on the cluster using the following command:

```
NCP> SET EXECUTOR ALIAS MAXIMUM LINKS n
```

- If the insufficient resources message occurred when the user was connecting to a nonalias node, increase the **MAXIMUM LINKS** value using the following command:

```
NCP> SET EXECUTOR MAXIMUM LINKS n
```

C. Do the following to check the page file quota value, and to increase it, if necessary:

1. Use the following command to display the process identification number (PID) for the **NETACP** process:

```
$ SHOW SYSTEM
```

2. Use the following command to display the current value for the page file quota value:

```
$ SHOW PROCESS/ID=netacp_pid/QUOTAS
```

3. If the page file quota is 0, increase the value specified for **NETACP\$PAGE_FILE**, using the following **DCL** command. Note that the default **NETACP\$PAGE_FILE** value is 8192.

```
$ DEFINE/SYSTEM NETACP$PAGE_FILE value
```

4. To cause the **NETACP\$PAGE_FILE** value to be permanently changed, modify the **SYSTARTUP.COM** file with the new value, and execute the **SYSTARTUP.COM** file before **STARTNET.COM** executes.

¹ The **MAXIMUM LINKS** value determines the maximum number of logical links permitted on a node simultaneously. You must consider the network configuration when determining an appropriate setting for the **MAXIMUM LINKS** parameter. However, a reasonable range for most networks is 25 to 50. The maximum value for **MAXIMUM LINKS** is 960. You must reduce this value to 512, however, if you also specify the **ALIAS MAXIMUM LINKS** parameter.

² The **ALIAS MAXIMUM LINKS** value determines the number of logical links permitted simultaneously on the cluster alias node. The maximum value for **ALIAS MAXIMUM LINKS** is 200. The default value is 32. If you specify **ALIAS MAXIMUM LINKS**, the maximum value permitted for the **MAXIMUM LINKS** parameter is reduced.

Insufficient Resources at Remote Node

5. Run NCP, and use the following command to shut down the network:

```
NCP> SET EXECUTOR STATE OFF
```

6. Use the following command to restart the network:

```
$ @STARTNET.COM
```

D. Do the following to check the MAXJOBS and MAXACCTJOBS values for the user who received the insufficient resources error, and to increase the values, if necessary and appropriate:

1. Run the Authorize Utility and use the following command to display information about the user's account:

```
UAF> SHOW user-id
```

2. Check the current values for MAXJOBS and MAXACCTJOBS.

MAXJOBS specifies the maximum number of batch, interactive, and detached processes that may be active at one time.

MAXACCTJOBS specifies the maximum number of batch, interactive, and detached processes that may be active at one time for all users who are on the same account as the specified user.

A value of 0 for MAXJOBS or MAXACCTJOBS indicates that an unlimited number of batch, interactive, and detached processes may be active at one time.

3. If an increase in the values is necessary or appropriate, use the following command to modify the values:

```
UAF> MODIFY user-id/MAXJOBS=n/MAXACCTJOBS=n
```

Recommendations

Some network servers (such as VTX and VAX Notes) can be heavily used. As a result, users who try to connect to these servers may encounter the insufficient resources message. You may want to allow more links to these servers by specifying a higher MAXIMUM LINKS value. However, specifying a higher MAXIMUM LINKS value can adversely affect performance, so weigh your decision to provide more links against the performance needs of the local users.

Invalid Parameter Value

Symptoms

While starting DECnet, the system displays the following message from NCP:

```
%NCP-W-INVPA, Invalid parameter value, Physical Ethernet address
Line = xxx-n
```

Explanation

This symptom indicates a DECnet-VAX node problem that results from protocols (such as LAT, DECElms, customer-written applications, or other Ethernet applications) starting before DECnet. Usually, this is because the LTLOAD.COM file is called before STARTNET.COM.

Troubleshooting Strategy

Make sure that DECnet starts first. To do this, stop all other protocols, restart DECnet, and then restart the other protocols.

Troubleshooting Procedure

1. Use the following command to see if other protocols are running, specifying the device type as one of the following:

Device-type designation	Device
XE	DEUNA, DELUA
ET	DEBNA
XQ	DEQNA, DELQA, DESQA
ES	DESVa

```
$ SHOW DEVICE device-type
```

The display from this command shows the device name and its current status.

2. Make sure that SYSSMANAGER:LTLOAD.COM is called from SYSTARTUP.COM, and is called after STARTNET.COM. To do this, make sure that the SYSSSYSTARTUP.COM file contains the following lines, in the following order:

```
@SYSSMANAGER: STARTNET.COM
@SYSSMANAGER: LTLOAD.COM
```

Or, you can submit a user-specified command file that contains the preceding commands in the correct order.

3. Run LATCP, and use the following command to stop the LAT protocol:

```
LCP> STOP NODE
```

Invalid Parameter Value

4. Use the following command to execute the STARTNET.COM file and restart the network:

```
$ @STARTNET.COM
```

5. Use the following command to execute the LTLOAD.COM file and restart the LAT protocol:

```
$ @LTLOAD.COM
```

LAN Bridge Cannot Downline Load

Symptoms

A LAN Bridge 100 or 150, intended to be used as a LAN Traffic Monitor, cannot downline load.

Explanation

This symptom indicates a LAN problem involving MOP protocol. The bridge has successfully completed the self test, but cannot downline load the LTM image due to a problem with the bridge setup, or with the load host.

Troubleshooting Strategy

1. Verify the setup of the bridge.
2. Check the load host for problems preventing it from downline loading the LTM software.
3. Use either DECelms or the switches on the bridge to correct the setup and enable the bridge to downline load. (Do not use both DECelms and the switches.)

Troubleshooting Procedure

1. Do one of the following to verify the hardware version:
 - Check the metal tag on the bridge for the hardware version.
 - Run DECelms, and use the following command to display the hardware version:

```
ELMS> USE bridge-id
ELMS> SHOW CHARACTERISTICS
```

To be able to downline load and function as a LAN Traffic Monitor, the bridge hardware must be at least Rev. E. The display shows various bridge characteristics including the ROM firmware version. A firmware version of 2.0 or greater equates to hardware Rev. E.

2. Use the following NCP command to see if circuit service is enabled on the host node:

```
NCP> SHOW CIRCUIT circuit-id CHARACTERISTICS
```

If circuit service is not enabled, use the following NCP commands to enable it:

```
NCP> SET CIRCUIT circuit-id STATE OFF
NCP> SET CIRCUIT circuit-id SERVICE ENABLED
NCP> SET CIRCUIT circuit-id STATE ON
```

LAN Bridge Cannot Downline Load

Note

The following command is optional. It enables service for the circuit in the NCP permanent database. However, you may not want to permanently enable service for the circuit due to the effect it has on performance.

```
NCP> DEFINE CIRCUIT circuit-id SERVICE ENABLED
```

3. Make sure that the cabling is connected securely.
4. Check the bridge indicator lights.

When the bridge is set up to function as a bridge, the indicator lights normally operate as follows:

- a. When you turn the bridge on, all the lights go on briefly, then all go out except the DC OK light.
- b. After about 15 seconds, the self test completes and the SELF TEST light goes on.
- c. After about 30 seconds, the ONLINE light goes on, unless the bridge is in a loop with another bridge or repeater connecting the two segments. In this case, the bridge may go into BACKUP state, and the ONLINE light may not come on.
- d. Finally, the activity lights begin blinking to indicate network activity. If the network is very busy, the lights blink very quickly and appear to be on continuously.

However, when the bridge is set up to function as a LAN Traffic Monitor, and the LTM software has been loaded, the ONLINE light blinks on and off in a pattern. This pattern indicates that the bridge is operating as a LAN Traffic Monitor.

5. If the bridge is intended to operate as a LAN Traffic Monitor, but the ONLINE light is not blinking, insert loopback connectors into the A and B ports, wait about 45 seconds, and check the lights again.

Table 5-3 shows the status of the bridge when various indicator lights are on.

Table 5-3 LAN Bridge 100 or 150 Indicator Lights

Indicator lights	Status
SELF TEST is off	The bridge has a hardware problem.
ONLINE, DC OK, and SELF TEST are on, and ACTIVITY lights are blinking approximately once per second	The bridge is set up to function as a bridge.
SELF TEST and DC OK are on, ONLINE is off, and the activity lights are blinking	The bridge is set up for downline loading and use as a LAN Traffic Monitor.

6. If the indicator lights show that the bridge is not set up properly for downline loading, use DECelms or the switches on the bridge to set up the bridge for downline loading.

LAN Bridge Cannot Downline Load

7. Make sure that the downline load switch (number 5) is disabled (UP).

Note

You can use either DECelms or the bridge switches. Do not use both for the same load. Step a, as follows, describes how to use DECelms to set up the bridge for downline loading. Step b describes how to use the bridge switches.

- a. To use DECelms to set up the bridge for downline loading, do the following:
 - i Use the following command to specify the environment for the remaining commands of this procedure:

```
ELMS> USE bridge-id
```
 - ii Use the following command to set the software downline load request flag:

```
ELMS> SET LOAD SWITCH TRUE
```

For a LAN bridge 150, you can also specify a password with this command.
 - iii Use the following command to specify the downline load file name:

```
ELMS> SET LOAD FILE "filename"
```

Note

The file name must be exactly 10 characters long.

For a LAN bridge 150, you can also specify a password with this command.

- iv Use the following command to cause DECelms to reset itself with the new information you specified:

```
ELMS> INIT
```
- b. To use the bridge switches to set up the bridge for downline loading, do the following:
 - i Clear the downline load switch and downline load information in NVRAM using the following steps:

Note

The following procedure sets *all* bridge parameters to the default settings. If you do not want to reset all the parameters, use DECelms instead.

- a. Press switch 2 (NVRAM RESET) down.
- b. Turn the bridge off.
- c. Turn the bridge on.
- d. When the self test completes, turn the bridge off.

LAN Bridge Cannot Downline Load

- e. Press switch 2 (NVRAM RESET) up.
- ii Make sure that the load host is properly set up, as described in the LTM installation documentation.

LAN Bridge Cannot Downline Load

- iii Connect at least one port to the same LAN or extended LAN as the load host. You can connect the other port to another LAN segment, or insert a loopback connector in it.
- iv Set the bridge switches as follows:

Note

For switches 3 and 4, you need not set both switches as long as you set the port switch that is in the same segment as the load host.

Table 5–4 shows the switch settings for the LAN Bridge 100 or 150.

Table 5–4 LAN Bridge 100 or 150 Switch Settings

Switch Number	Name	Setting
1	Manufacturing Mode	Up (Off)
2	NVRAM Reset	Up (Off)
3	Port A Access	Down (On)
4	Port B Access	Down (On)
5	Downline Load	Down (On)
6	Not Used	Up (Off)

- v Turn the bridge on.
The load takes less than five minutes, unless the load host is extremely busy.

LAN Segment Communication Problem

LAN Segment Communication Problem

Symptoms

All the systems on a LAN segment are unable to communicate with systems beyond their segment. However, all the systems on the isolated LAN segment can still communicate among themselves.

Explanation

This symptom indicates a LAN problem involving the physical layer. The symptoms could be related to problems causing a bridge or repeater to segment.

Bridges connect Ethernet LANs to create extended LANs, and repeaters connect Ethernet segments to expand a LAN. Bridges keep the traffic between systems on a LAN segment within that LAN segment, and out of the general network traffic on the extended LAN. Restricting network traffic this way keeps segment traffic to a minimum and prevents unnecessary traffic from entering the extended LAN. Repeaters do not isolate traffic; however, if a repeater detects faulty signals that cause a high number of collisions, the repeater automatically stops repeating the signals until it detects good signals again.

Potential causes of this problem include the following:

- Occasionally, a repeater or bridge may fail, or may be disconnected accidentally, causing an entire segment to become isolated from the rest of the extended LAN.
- A problem may exist on the LAN on the other side of the device (for example, a babbling device) that causes the bridge or repeater to segment.
- A faulty H4000 tap for the bridge or repeater may cause the device to segment.

Troubleshooting Strategy

To begin solving this problem, use your knowledge of the network topology, and your network map to isolate the source of the problem to the interconnecting device for the isolated LAN segment. After you determine the device that is causing the problem, continue with the following steps:

- For bridge problems at sites where DECelms is available, use step A.
- For bridge problems at sites where DECelms is *not* available, use step B.
- For bridge problems at sites using ETHERnim, use step C.
- For repeaters (including DEREPE and DEREN (Ethernet Repeaters), DEMPR (ThinWire multipoint repeater), and DESPR (ThinWire singlepoint repeater), use step D.
- For H4000 problems, use step E.

Troubleshooting Procedure

A. If DECelms is available at your site, do the following to check the bridge:

1. Run DECelms, and use the following command to verify that the bridge lines are operating and are in the FORWARDING state:

```
ELMS> USE bridge_id  
ELMS> SHOW KNOWN LINES STATUS
```

The display shows line characteristics for the lines on the bridge, and whether the lines are in the forwarding state.

2. Use the following command to display the bridge counters:

```
ELMS> SHOW COUNTERS
```

The value for the bridge seconds counter tells you how long the bridge has been running.

3. At sites where several users have access to the DECelms software, another user may have mistakenly set up the bridge to filter all packets destined for certain addresses. If you suspect this is the case, use the following command to verify the forwarding database status on the bridge:

```
ELMS> SHOW ADDRESS address
```

The display shows the forwarding entry for the address specified.

4. Check the display for the destination address.

The designation NONE means that the bridge does not forward packets to that address. If the designation is NONE, do the following:

- a. Make sure that the bridge should be forwarding packets to that address.

Sometimes packet forwarding is intentionally disabled for a particular address.

- b. If packets should be forwarded to the address, use the following command to enable the bridge to forward packets to the address:

```
ELMS> REMOVE ADDRESS address PASSWORD password
```

In the normal course of operation, the bridge learns the correct action to take for packets destined for this address.

B. If DECelms is not available at your site, do the following to check whether the bridge is off line for segmentation:

1. Go to the bridge for the isolated LAN segment.
2. Make sure that the power is on.
3. Make sure that the cable connecting the bridge to the H4000 transceiver or DELNI is properly connected.

LAN Segment Communication Problem

4. Check the indicator lights on the bridge.

Normally, the activity lights blink on and off for each packet sent. If the bridge is processing many packets, the lights are on continuously. This is not unusual, and does not indicate a problem. However, if the lights are off, there is probably a hardware or power problem with the bridge.

5. If the activity lights are off, replace the bridge.

6. If the ONLINE light is off, try to connect with a network node on the other side of the bridge to make sure the bridge is not in a loop configuration with another bridge or repeater (and is in a backup state). If the connection succeeds, there is a loop, and this bridge is in backup mode. In this case, the loop is not a problem. If a failure occurs on the network, this bridge changes from backup to online mode.

7. If no unexpected loops exist and the problem persists, replace the problem bridge with another bridge to verify if the problem is hardware related. If the new bridge functions properly, the problem with the old bridge is probably hardware-related.

- C. **If you have ETHERnim at your site, use it to poll or show nodes on the segment in question, and on each previous segment, until you locate the source of the problem.**

- D. **Use the following procedure to solve problems relating to repeaters:**

1. Go to the repeater for the isolated LAN segment.
2. Make sure that the power is on.
3. Make sure that the cable connecting the repeater to the H4000 transceiver or DELNI is properly inserted.
4. Run the self-test.
5. Check the indicator lights on the repeater.
The SEGMENTED light usually indicates a circuit problem.
6. See the appropriate repeater manual for further corrective actions.

- E. **If the bridge is still not reachable, the H4000 transceiver that connects the bridge to the local segment may be faulty. To determine if this is a problem, do the following:**

1. Go to a node on the other segment to see if that node can communicate through the bridge.

If the node on the other segment can reach the bridge, then the H4000 transceiver that connects the bridge to the local segment is probably broken.

LAN Segment Communication Problem

2. Check that connections are secure.
3. If the connections are secure and the bridge is still unreachable, move the bridge to another H4000 tap or DELNI port.

Recommendations

- If possible, include redundant bridges and repeaters on your network. The redundant devices provide service if a primary device fails, helping to ensure uninterrupted service for your users.
Digital's bridges and repeaters perform automatic failover when the network includes redundant devices. To provide backup service for bridges on your network, you need only provide one bridge to act as the backup for all bridges on your network, because bridges use a spanning tree algorithm.
- If you are using the bridge as a LAN Traffic Monitor, and are not using both of the bridge's ports (A and B), insert a loopback connector in the unused port. If you do not use the loopback connector, the bridge cannot complete the self-test when it starts running. In the event of a bridge failure elsewhere on the network, you can temporarily use the LAN Traffic Monitor bridge as a backup.
- Be sure to follow the guidelines for configuring your network with bridges. The maximum number of bridges permitted in a linear setup (from point A to point B) is seven.
- Be sure to follow the guidelines for configuring repeaters on your network. The maximum number of repeaters permitted in a linear setup (from point A to point B) is two. (A pair of fiber-optic repeaters count as one repeater.)
- Keep track of when and where new taps are installed on the network.

LAT Port Hung

LAT Port Hung

Symptoms

A user's terminal does not respond.

Explanation

This symptom indicates a LAN problem, and can result from any of the following:

- User's terminal is defective
- Wiring is disconnected
- LAT port is hung
- Setup information is not consistent between the LAT port and the terminal

Troubleshooting Strategy

To resolve this problem, first isolate the source of the problem to the wire, terminal, or LAT port. After you have determined the source of the problem, resolve it as appropriate to the source of the problem.

Caution

The following procedures involve initializing a LAT device and logging out a LAT port. Both of these procedures can be disruptive to network users.

Initializing a LAT device causes the LAT device to reload, disconnecting all sessions on all ports from and to the LAT device. As a result, you may want to initialize devices during off hours using the /AFTER command qualifier.

Logging out a LAT port disconnects all sessions to that port, and may cause data to be lost for the user on that port.

Troubleshooting Procedure

1. To ensure the problem is not due to user error, do the following at the terminal in question to verify that the user's terminal does not respond:
 - a. Press Ctrl/Q to send an X/ON character through the network.
 - b. Press the Break key.
 - c. Press Return.
2. Check to see if the terminal setup (speed, parity, stop bits, character size, modem control, and so forth) is correct.
3. Trace the wire or check the network map to determine the LAT device and port to which this terminal is connected.
4. Make sure that the LAT port setup matches the terminal setup.
5. Use one of the following procedures to test the terminal server port, the user's terminal, and the line between the terminal server and the terminal.

Your choice depends on whether or not your environment is a DECconnect environment.

With DECconnect, you can swap and test ports at the satellite equipment room (SER) or office communications cabinet (OCC) using simple patching mechanisms. You can also use these techniques in a non-DECconnect environment, but without the ease of using the patch panels.

For DECconnect environments, do the following:

- a. Find the DECserver port in the SER or OCC, and swap the user's port with another port that you know is working.
- b. Test the port at the user's office. If the new port works, the user's old port is faulty. Call Customer Services.
- c. If the preceding steps have not resolved the problem, use a terminal in the SER or OCC on the user's port.

If this terminal works, then you know the terminal server is working properly.

- d. Do one of the following to determine if the problem is the terminal or the line between the terminal server and the terminal:
 - Try a different terminal on the user's LAT port.
If the terminal works, then the user's terminal is faulty. If the terminal does not work, then the line between the terminal server and the terminal is faulty. Call Customer Services.
 - Try the user's terminal on a different LAT port.
If the terminal does not work, then the terminal is faulty. If the terminal works, then the line between the terminal and the terminal server is faulty. Call Customer Services.
- e. Log the port out to clear the port.
- f. Initialize the DECserver at a later time.

In a non-DECconnect environment, do the following:

- a. Connect to the server in question, using TSM *or* NCP as follows:

```
TSM> USE SERVER server-id
```

```
NCP> CONNECT NODE server-id
```

- b. Use the TEST PORT commands to test the port internally.
If this test fails, there is a hardware failure on the LAT device.
- c. Check the server and port for errors.
- d. Log the port out.
- e. Initialize the LAT device at a later time.

LAT Port Hung

- f. If logging the port out does not clear the port, the LAT device or card probably needs to be replaced.
- g. Do an external loopback test with a loopback connector.

This test can be performed anywhere along the path between the LAT device and the terminal to test the LAT device and the portion of the line that is looped back.

 - i Try looping at the back of the LAT device.

If this fails, then the problem is in the LAT device hardware.
 - ii Try looping back at the user's terminal.

If the test fails at the terminal, then the line between the LAT device and terminal is faulty and needs to be repaired.

If this test is successful, and the terminal still does not work, then the terminal is probably defective.
 - iii To confirm whether the terminal is defective, try using the terminal on a line you know is working, or try a terminal you know is working on the faulty line.

LAT Print Queue Problems

Symptoms

A LAT print queue on a VMS system is not printing or is printing incorrectly. The print queue may be in a stopped, stalled, or paused state, or jobs may be retained on error.

Note

The troubleshooting procedures for this problem apply only to LAT print queue problems on VMS hosts.

Explanation

This is a LAN problem involving the LAT protocol, resulting from any of the following causes:

- VMS host system setup problems
 - Print queue setup (for example, the wrong queue processor)
 - Device setup (for example, improper LTA device setup resulting in the LTA device pointing to the wrong server)
 - Mismatch between VMS host system definitions and terminal server definitions
- Terminal server setup problems
 - Incorrect port and server definitions
 - Multiple servers with the same name
- Printer problems
 - Incorrect setup
 - Hardware problems
 - Out of paper
- Ethernet transmission errors

Troubleshooting Strategy

Check the state of the queue on the VMS host system, and resolve any problems based on the print queue state. Ensure that parameters on the terminal server, VMS host system, and printer match appropriately, as follows:

- Server name on the terminal server matches the server name on the LTA device on the VMS host system
- Port name on the terminal server matches the port name on the LTA device on the VMS host system
- No duplicate server names exist

LAT Print Queue Problems

- Printer characteristics on the printer match the printer characteristics on the terminal server port

Note

To correct the LAT print queue problem, you can use TSM or NCP to enter commands on the terminal server, or you can go directly to the terminal server and enter the commands.

Whether you use TSM or NCP, or enter the commands directly on the server, the procedure is the same, only the prompts are different. The following steps show how to log in to the terminal server using either TSM or NCP.

1. To use TSM to log in to the server, run TSM and use the following command:

```
TSM> USE SERVER server-id
```

2. To use NCP to log in to the server, the server must be defined in the NCP database. If the server is defined in the NCP database, run NCP and use the following command to log in to it:

```
NCP> CONNECT NODE server-id
```

If the server is not defined, you can use the following NCP command:

```
NCP> CONNECT VIA service-circuit PHYSICAL ADDRESS-  
_NCP> ethernet-physical-address  
Console connected press cntrl D when finished.  
<carriage return>  
# <login password>
```

When you enter the following server commands, the system displays the Local> prompt rather than the TSM> prompt. To return to NCP from the Local> prompt, press Ctrl/D.

3. To enter commands directly on the server, go to the server and log in using your user name.

```
Enter username> username
```

When you enter the following server commands, the system displays the Local> prompt rather than the TSM> prompt.

Note

From this point on, the terminal server commands are the same regardless of how you access the server. However, the following procedures assume you use TSM to access the server.

Troubleshooting Procedure

Before you begin to solve this problem, do the following:

- Enable the display of all LAT error messages on your terminal using the following command:

```
$ SET MESSAGE SYS$MESSAGE:NETWRKMSG
```

To enable display of LAT error messages for all users, put the SET MESSAGE command in the SYS\$MANAGER:SYLOGIN.COM file.

- Determine the state of the queue on the VMS host system using the following command.

```
$ SHOW QUEUE queue_name/FULL
```

Table 5–5 shows the queue states and conditions, their meanings, and the troubleshooting step that solves the problem.

Table 5–5 Print Queue States and Conditions

State	Meaning	Solution
Stopped	Print queue has been stopped using the STOP/QUEUE /RESET command, or the queue manager has been stopped, or some other problem has occurred such as termination of the symbiont process.	Step A
Stalled	Printer is unable to complete the request due to a flow control problem, such as lack of paper.	Step B
Paused	Print queue has experienced an error or unexpected event when communicating with the server.	Step C
Printing Incorrectly	Printer is generating incorrect output.	Step D
Retained on Error	A job experienced an error during execution, but remains in the queue.	Step E

A. For a stopped queue, start the queue on the VMS host system using the following DCL command:

```
$ START/QUEUE queue-name
```

A stopped queue usually indicates a problem on the VMS host system.

1. Check SYSSYSTEM for LATSYS.DMP files.
2. If LATSYS.DMP files exist, then check the OPERATOR.LOG file for jbc errors that explain why the files exist.

For further help with LATSYS problems, call your local Digital customer service representative.

B. For a stalled queue, do the following:

1. Use the following TSM command to get information about the printer:

```
TSM> USE SERVER server-name
TSM> SHOW PORT port-id STATUS
```

- a. If the "Status" field indicates that your node is not connected to the terminal server, check the queue on the terminal server, using the following commands:

```
TSM> USE SERVER server-name
TSM> SHOW QUEUE
```

If the display shows entries in the queue for the VMS host system, and the port is connected to another system, then the server is waiting for the port to become idle.

Also, font loading in some types of printers can take time. The printer is not printing, check to see if the printer is loading fonts before you take any other action.

LAT Print Queue Problems

- b. If signal check is enabled on the port, and the port status is "signal wait," check the "Input Signals". If they show no modem signals from the printer, then the server is waiting for a modem signal from the printer.
 - Make sure that no problem exists on the printer (such as a broken or disconnected cable) that might prevent the printer from sending signals.
 - After you resolve any problems that might prevent the printer from sending signals, turn the printer off, then on, to clear the problem.
 - c. If the flow control fields ("Input" and "Output") indicate that output is XOFFed, then the server is waiting for an XON character from the printer. The printer may have sent an XON character and the server lost it, or the printer may have never sent an XON character.
 - Make sure that the printer is functioning properly. For example, make sure it is on, has paper, and is not displaying hardware errors.
 - Turn the printer off, then on, to clear the problem.
 - As a last resort, log out the port.
2. If the printer is not printing after a reasonable amount of time, and is not loading fonts, use the following command to log out the terminal server port and clear the problem:

```
TSM> LOGOUT PORT port-number
```

C. For a paused queue, do the following:

1. Use the following DCL command on the VMS host system to display the LAT device identification:

```
$ SHOW QUEUE queue-name/FULL
```
2. Use the following DCL command on the VMS host system to unspool the LAT device:

```
$ SET DEVICE /NOSPOOL LTAn:
```
3. Copy a printable file to the LAT device to confirm whether the LAT device is working:

```
$ COPY/LOG file_name LTAn:
```

If the COPY command succeeds, go to step 8.

If the COPY command does not work you get the error, "data set hangup." The "data set hangup" error indicates that the VMS host cannot find the terminal server. This problem is due to any of the following:

- Incorrect or mismatched information on the terminal server and the VMS host
- Data transmission problem on the Ethernet segment

- Port configuration problems
 - Terminal server problems
4. Make sure the server name and port name on the VMS host system are correct and match the server name and port name on the terminal server.

- a. Do the following on the VMS host system to display the port characteristics:

```
$ MCR LATCP
LCP> SHOW PORT LTAn:
LCP> EXIT
```

- b. Do the following on the terminal server to display the server and port characteristics:

```
$ TSM
TSM> SHOW SERVER
TSM> SHOW PORT xx
TSM> EXIT
```

- c. If the host system definitions are wrong, use the following LATCP command to redefine them. Make sure that the LTA port on the VMS host system is defined to be an applications port, and that the QUEUE attribute is set.

```
LCP> SET PORT LTAn:/APPLICATION/NODE=servername/PORT=portname
```

- d. If the server characteristics on the terminal server are wrong, use the following commands to redefine them.

```
TSM> SET SERVER server-id characteristics
TSM> DEFINE SERVER server-id characteristics
```

Note

The SET command causes the change to take effect immediately. The DEFINE command makes the change permanent, so that when you reboot the server, the new characteristics are in place.

The DECserver 500 and DECserver 550 do not use a DEFINE command. Instead, for these terminal servers, use the SET command to make the changes. To cause the changes to take effect permanently, modify the load image.

See the DECserver 500/550 documentation for more information on modifying a load image and configuring a DECserver 500 terminal server.

- e. If the port characteristics on the terminal server are wrong, use the following command to redefine them:

```
TSM> DEFINE PORT port-id characteristics
```

- f. Log the port out to cause the port changes to take effect.

```
TSM> LOGOUT PORT port-number
```

LAT Print Queue Problems

5. Check to see if data transmission problems exist on the Ethernet segment.

- a. Run TSM and use the following command to display the terminal server counters:

```
TSM> SHOW SERVER COUNTERS
```

- b. Check the "Solicitations Accepted" and "Solicitations Rejected" fields.
- c. Try the COPY command again.
- d. Display the terminal server counters again and check the "Solicitations Accepted" and "Solicitations Rejected" fields.
 - If the "Solicitations Rejected" field increments, then the server is operating but is not accepting solicitations, probably because the port is misconfigured. Go to step 6.
 - If neither field increments when you issue the COPY command, the terminal server did not receive the copied file. A cabling or network problem could exist. Go to step 7.
 - If the "Solicitations Accepted" field increments, the COPY command worked. Go to step 8.

6. If the "Solicitations Rejected" field increments, make sure the terminal server port is properly configured.

- a. Run TSM, and use the following command to display the current characteristics for the server port:

```
TSM> SHOW PORT
```

The terminal server port for the print queue must have the following characteristics defined:

- Unique name
 - Access remote
 - Autobaud disabled
 - If group codes are enabled, the port's group code must match the VMS host group code definitions
- b. If the terminal server characteristics are not correct, use the following TSM command to specify the correct characteristics, as required:

```
TSM> DEFINE PORT characteristics
```

- c. Log the port out to cause the changes to take effect.

```
TSM> LOGOUT PORT port-number
```

- d. Check the printer to make sure that the printer's hardware is functional. For example, make sure that the printer is connected, is on, has paper, and so forth. Continue with step 9.

7. If neither the "Solicitations Accepted" nor the "Solicitations Rejected" field increments, check to see if cabling problems exist between the terminal server and the VMS host, or if a network problem exists.

- a. Use the following command to see if the terminal server knows about the VMS host system:

```
TSM> SHOW NODE node-id
```

- b. If the terminal server displays the VMS host system, then the COPY command should work.
- c. If the terminal server does not display the VMS host system, make sure that the terminal server cabling is intact and properly connected, and that the VMS host system cabling is properly connected to the Ethernet.
- d. Use NMCC/VAX ETHERnim to verify the path to the terminal server and the VMS host system.

Use LAN Traffic Monitor data to see if a LAN segment problem or babbling device problem exists on the network. (See the procedures for "LAN Segment Communication Problem" and "Babbling Device" in this chapter.)

Use NCP loopback commands or NMCC/VAX ETHERnim to isolate network problems.

- e. Try to reboot the terminal server to clear the problem.

Note

Rebooting the terminal server disconnects all connections to the terminal server. Use this step only as a last resort.

8. If the "Solicitations Accepted" field increments, and the COPY command succeeds, but you still cannot print using the print queue, the queue could be using the wrong queue processor.

Use the following DCL command to initialize the queue and define the correct queue processor:

```
$ INITIALIZE QUEUE queue_name/PROCESSOR=LATSYM
```

9. Confirm that the port is working properly using the following command:

```
TSM> TEST PORT
```

The TEST PORT command allows the terminal server to communicate directly with the printer without VMS. It helps confirm that the connection between the terminal server and the printer is working.

- If the test pattern prints correctly, the port is operating properly.
- If the test pattern does not print correctly, check the connection between the terminal server and the printer. Also confirm that the characteristics of the terminal server and printer match, and that the printer is operating properly, as shown in step 6a.

D. For a queue that is printing incorrectly, do the following:

LAT Print Queue Problems

Make sure that the following parameters are the same on *both* the printer and the terminal server port:

- Port speed
- Flow control
- Character size
- Parity
- Baud rate
- Autobaud is disabled on the terminal port

1. Use the following command to check the terminal server port:

```
TSM> SHOW PORT port-id CHARACTERISTICS
```

2. If you need to change the terminal server characteristics, use the following command:

```
TSM> DEFINE PORT port-id characteristics
```

3. Log the port out to cause the changes to take effect.

```
TSM> LOGOUT PORT port-number
```

4. For instructions on setting printer characteristics, see the printer hardware documentation.

E. For a queue displaying the "retained on error" message, do the following:

1. Use the following command to display any additional messages:

```
$ SHOW QUEUE queue_name/FULL
```

Usually, the "retained on error" message results in a stalled or paused queue. Jobs that abort and are retained on error are usually the result of one of the following reasons:

- Port is logged out in the middle of a job
- LAT shuts down on the VMS host system
- Ethernet device fails while transmitting
- User deletes the job from the queue using the DELETE/ENTRY=nnn command

LAT Print Queue Problems

2. If the queue is stalled or paused, use the steps for solving stalled or paused queues.
3. After you solve the stalled or paused problem, use the following DCL command to clear the queue of the job retained:

```
$ SET ENTRY/RELEASE job_entry_number
```

Recommendations

For more information about print queue operations, see the *Guide to Maintaining a VMS system*.

Line Synchronization Lost

Line Synchronization Lost

Symptoms

A circuit goes down and up every two or three seconds, and the system displays the following DECnet event messages:

```
%%%%%%%%%% OPCOM 27-JUN-1988 14:22:06.17 %%%%%%%%%%%  
Message from user DECNET on NODE1  
DECnet event 4.7, circuit down, circuit fault  
From node x.xxx (NODE2), 27-JUN-1988 14:17:58.10  
Circuit DMC-3, Line synchronization lost
```

```
%%%%%%%%%% OPCOM 27-JUN-1988 14:22:06.17 %%%%%%%%%%%  
Message from user DECNET on NODE1  
DECnet event 4.10, circuit up  
From node x.xxx (NODE2), 27-JUN-1988 14:18:01.79  
Circuit DMC-3, Adjacent node = x.xxx (NODE3)
```

Explanation

This symptom indicates a DECnet-VAX node problem in which the data link protocols between the two nodes cannot be initialized. This symptom usually indicates a hardware or line problem, such as the following:

- Transceiver cable that is not properly connected
- Faulty communications board
- Improper system parameter settings for IRPCOUNT, LRPCOUNT, and SRPCOUNT

It can also indicate the following:

- Local Area VAXcluster set up improperly as a boot node
- Synchronous line problems (for example, modem or digital service unit problems)
- Faulty H4000 connection to the Ethernet

Troubleshooting Strategy

To solve this problem, determine if the problem is on the local node or the Ethernet, then complete the appropriate actions below:

- A. If the problem is on the local node, do the following on the local node:
 1. Make sure that the current values have not reached or exceeded the initially allocated values for IRPCOUNT, LRPCOUNT, SRPCOUNT.
 2. Check the circuit counters.
 3. Make sure that the cable connections are secure.
- B. If the problem is on the Ethernet, do the following on the local node:
 1. If the node is set up as a Local Area VAXcluster boot node, make sure that the VAXCLUSTER parameter in SYSGEN is set properly.
 2. Make sure that the current counter values have not reached the maximum permitted values.
 3. Check the circuit counters.

4. Check the line counters for open or short circuits on transmit and receive.
5. Make sure that the cable connections are secure.
6. Make sure that the Ethernet controller module is properly seated.
7. Use loopback tests as necessary, if the problem relates to synchronous devices and modems.

Troubleshooting Procedure

To begin solving this problem, check the events displayed on several nodes on the Ethernet to determine if the problem is related to the local node or the Ethernet. If the line synchronization lost message is displayed only on one node, the problem is on that node. If the line synchronization lost message is displayed on multiple nodes, the problem is on the Ethernet.

A. Use the following steps to resolve problems on the local node:

1. Use the following command to display the current values for IRPCOUNT (I/O request packet count), LRPCOUNT (large request packet count), and SRPCOUNT (small request packet count):

```
$ SHOW MEMORY/POOL/FULL
```

IRP, LPR, and SRP are three preallocated memory pools in the nonpaged pool area. The nonpaged pool area is a portion of physical memory permanently allocated to the system for the storage of data structures and device drivers. Its initial size is determined by AUTOGEN, but automatic expansion of the area occurs if necessary.

2. Compare the current values of IRPCOUNT, LRPCOUNT, and SRPCOUNT to the initial allocation.
3. If the current values equal or exceed the initial allocation, edit MODPARAMS.DAT to increase both IRPCOUNT, LRPCOUNT, and SRPCOUNT values, as well as IRPCOUNTV, LRPCOUNTV, and SRPCOUNTV values.

(IRPCOUNTV, LRPCOUNTV, and SRPCOUNTV values are the upper limits to which the IRPCOUNT, LRPCOUNT, and SRPCOUNT values can be automatically increased by the system.)

In determining the amount to increase the values, you must trade off the permanent allocation of memory for nonpaged pool against the small amount of CPU overhead required to do pool expansion. If physical memory on your system is limited, it may be reasonable to accept a low to moderate amount of expansion.

4. Execute the AUTOGEN.COM file to cause the changes to take effect.

Note

The AUTOGEN.COM command procedure reboots the system. Be sure you really want to reboot the system at this time before you execute the following command:

```
$ @SYS$UPDATE:AUTOGEN GETDATA NOFEEDBACK REBOOT
```

Line Synchronization Lost

5. For point-to-point circuits, the problem usually involves a cable that is not connected properly, a faulty communications board, or a faulty or noisy circuit. Run NCP, and use the following command to check the counters on the failing circuit for any errors:

```
NCP> SHOW CIRCUIT circuit-id COUNTERS
```

6. If the counters display the greater than symbol (>), then the counters have reached their maximums and cannot record any further changes. In this case, zero the counters.

```
NCP> ZERO CIRCUIT circuit-id
```

7. After a short time, check the counters again to see whether there is a change, and follow up on any unusual counter changes, such as high error rates.
8. If the counter information does not help determine the problem, use loopback tests (as described in NCP Loopback Tests).
9. Make sure that the cable connections are secure.

B. Use the following steps to resolve problems on the Ethernet:

1. Perform all the steps in A, then continue with the following steps.
2. If the node is the boot node for a Local Area VAXcluster, make sure that the SYSGEN parameter, VAXCLUSTER, is set to 1, and execute the AUTOGEN.COM file to cause the changes to take effect.

The VAXcluster parameter controls loading of the cluster code. The default setting is 1, which means to load if SCSLOA is being loaded. A setting of 0 means to never load. A setting of 2 means to always load, and always load SCSLOA.

Note

The AUTOGEN.COM command procedure reboots the system. Be sure you really want to reboot the system at this time before you execute the following command:

```
$ @SYS$UPDATE:AUTOGEN GETDATA NOFEEDBACK REBOOT
```

3. For Ethernet circuits, if the circuit down counter has incremented, the problem is due to a faulty hardware device, improperly terminated cables, or loose cable connections. In particular, the problem may be related to open or short circuits. Use the following command to check the counters on the line:

```
NCP> SHOW LINE line_id COUNTERS
```

Check the transmit and receive counters for open or short circuits. On Ethernet circuits, if there are no open or short circuits, the problem is probably due to a faulty communications board.

4. Make sure that the cable connections are secure.
5. Connect the transceiver cable to another H4000 or DELNI to determine if the problem is related to an H4000 transceiver failure or to a faulty H4000 connection to the Ethernet.

Line Synchronization Lost

If the new connection works, then the problem was related to the H4000 tap into the Ethernet, or to the H4000 transceiver itself.

6. To resolve H4000 problems, first try to retap the H4000 into the Ethernet. If the problem persists after retapping the H4000, the H4000 may be faulty.
7. Replace the H4000, and check the DECnet event messages again.
8. For problems relating to synchronous devices and modems, use loopback tests. See NCP Loopback Tests for more information on these tests.
9. For controller or device level problems, ask Customer Services to make sure that the Ethernet controller module is seated properly.

Login Incorrect

Login Incorrect

Symptoms

A user on an ULTRIX system receives the following message when attempting to access a remote host:

```
Login incorrect
```

Explanation

This symptom indicates an ULTRIX host problem involving the internet protocol (IP). The user specified an incorrect account or password, or both, when attempting to access a remote host.

Note

This symptom may not indicate a problem. It is possible that the user is not intended to have an account on the remote host. Before you try to resolve this problem, be sure the user is intended to have access to the remote host.

Troubleshooting Strategy

1. If the user is intended to have access to the remote host, but does not have an account on the remote host, log in to the superuser account and create an account for the user.
2. If the user has an account on the remote node but cannot access it due to problems with the password, log in to the superuser account and modify the user's password in the `/etc/passwd` file.

Troubleshooting Procedure

1. If access is determined according to the `/etc/passwd` file, use the following command to display the contents of the `/etc/passwd` file, where `username` is the user's login name.

```
# grep username /etc/passwd
```

If YP is the method used to determine access, use the following command instead:

```
# ypcat passwd | grep username
```
2. Look for an entry for the user in the `/etc/passwd` file. If no entry exists for the user, go to step a. If an entry exists, but is incorrect, go to step b.
 - a. If no account exists for the user, and the user is intended to have access to the remote host, execute the following command from the superuser account:

```
# adduser
```

When you execute the `adduser` command, the system displays questions you must answer regarding the account you are creating. Answer the questions appropriately for the user.

- b. If an account exists, but the user cannot recall the password, use the following command to define a new password for the user

```
# passwd username
```

Login Information Invalid

Login Information Invalid

Symptoms

Users receive an error message such as the following when attempting any network operation except SET HOST:

```
%MAIL-E-LOGLINK, Error creating network link to node NODEID  
-SYSTEM-F-INVLOGIN, login information invalid at remote node
```

Explanation

This symptom indicates a DECnet-VAX node problem involving the session layer. It can be caused by any of the following:

- User-supplied access control information is incorrect.
- Proxy access is set up incorrectly.
- The user and password for a specific object on the remote node does not match a valid account in the System User Authorization (SYSUAF) file.
- The nonprivileged password defined in the executor characteristics on the remote node does not match the password defined in the remote node's SYSUAF file.
- The executor does not have a nonprivileged user or nonprivileged password defined.
- AUTHORIZE parameter settings for the default DECnet account may cause this message. For example, this message can occur if the DISUSER flag is set or the account is expired.

If the login information invalid error occurs intermittently, the remote system is probably a cluster system that has a node or nodes set up improperly. When the login information goes to the improperly set up node, the error message occurs. However, if the login information goes to a properly set up node, the login is successful.

To troubleshoot this problem effectively, you need to understand the order of access control for VMS systems. VMS systems permit access based on the type of access control information the system receives.

VMS systems check first for user-supplied access control information. If user-supplied information does not exist, the VMS system checks for proxy access information. If proxy access information does not exist, the VMS system checks for default access information.

Figure 5-2 shows the order of access control for VMS systems in more detail.

Figure 5–2 VMS Access Control

Note

NCP parameters determine the type of proxy access permitted, if any, for objects on a node or for the node itself. NCP information specified for an

Login Information Invalid

object always supersedes the executor information. This is true for proxy as well as default account information.

Proxy access for objects and the executor is determined according to the parameters shown in Table 5–6.

Table 5–6 NCP Proxy Access Parameters

Object Parameter	Function
PROXY INCOMING	Allows proxy login to the object.
PROXY OUTGOING	Allows the object to initiate proxy login.
PROXY BOTH	Allows both incoming and outgoing proxy login access. This is the default.
PROXY NONE	Prohibits incoming and outgoing proxy login access. If you omit the PROXY parameter, proxy access is determined according to the executor parameters.

Executor Parameter	Function
INCOMING PROXY DISABLED	Ignores all incoming proxy requests, and instead, relies exclusively on access control information supplied in the connect requests to validate the logical link.
INCOMING PROXY ENABLED	Invokes the appropriate proxy, based on the source user, source node, and supplied access control information, if any. This is the default.
OUTGOING PROXY DISABLED	Specifies that proxy login is not requested on any outgoing logical links.
OUTGOING PROXY ENABLED	Specifies that proxy login is requested on outgoing logical links. This is the default.

Troubleshooting Strategy

Before you attempt to solve this problem, do the following:

- Determine the type of access control the user specified when trying to access the remote node.
 - Determine the objects the user tried to access.
 - If the error occurs intermittently, the remote node may be a cluster system. Determine which node of the cluster is improperly set up, and correct the set up.
- A. For user-specified access control problems, do the following:
1. Check the user-specified access control information.
 2. Make sure that the account the user wants to access exists on the remote node.
 3. Modify the password information on the remote account.

- B. For proxy access control problems, do the following on the source (local) and target (remote) nodes:
 - 1. Do the following on the source (local) node:
 - a. Make sure that PROXY OUTGOING or PROXY BOTH is enabled for the object.
 - b. Make sure that OUTGOING PROXY is enabled for the executor.
 - 2. Do the following on the target (remote) node:
 - a. Make sure that PROXY INCOMING or PROXY BOTH is enabled for the object.
 - b. Make sure that INCOMING PROXY is enabled for the executor.
 - c. Make sure that NETPROXY.DAT has the correct proxy definitions.
 - d. Make sure that the proxy account exists in the SYSUAF file.
- C. For default access control problems, do the following:
 - 1. If the requested object has a user name and password associated with it, make sure that the definitions in the SYSUAF file match those specified in NCP.
 - 2. If the requested object does not have a user name and password associated with it, check to see if the remote node's executor has a nonprivileged user and nonprivileged password specified.
If they are specified, make sure that the definitions in the SYSUAF file match those specified in NCP.
 - 3. If neither the object nor the executor have a nonprivileged user name and password, define them as necessary.

Troubleshooting Procedure

Before you try to solve this problem, do the following:

- 1. Use the following list to help determine which step to use to solve the problem:
- 2. If the user specified explicit access control information when trying to access the remote node, use step A to solve the problem.
- 3. If the user tried to access the remote node using proxy access, use step B to solve the problem.
- 4. If the user did not try to access the remote node using proxy access, use step C.
- 5. If you do not know if the user tried to access the remote account using proxy access, use step B first, and continue with step C if necessary to solve the problem.
- 6. If the error has been occurring intermittently, the remote node may be a cluster system. Determine which node of the cluster is improperly set up, and correct the set up.

Note

To help solve this problem faster, you can try to connect to a different object on the remote node. If the connection succeeds when directed

Login Information Invalid

toward the new object, then the problem is probably object-specific, and you can focus your efforts on ensuring that the object on the remote node is set up properly.

A. For explicit access control problems, do the following:

1. Try to log in to the remote account that the user tried to access, using the access information the user specified.

If you cannot log in, the access information is probably incorrect.

2. Log in to an account on the remote node that has SYSNAM and SYSPRV privileges, and perform the following steps.

3. Use the following AUTHORIZE command to check that the user has an account in the remote node's system user authorization file (SYSUAF):

```
UAF> SHOW user-id
```

4. Use the following AUTHORIZE command to define a new password for the user's account:

```
UAF> MODIFY user-id/PASSWORD=password
```

5. Try logging in to the user's account with the new password.

B. For proxy access control problems, do the following on the source (local) and target (remote) nodes:

1. Do the following on the source (local) node:

- a. Run NCP, and use the following command to display the current settings for proxy access to the object:

```
NCP> SHOW OBJECT object-name CHARACTERISTICS
```

If PROXY OUTGOING is not specified, and if proxy access is required for the requested object, use the following NCP command to enable OUTGOING proxy access for this object only:

```
NCP> SET OBJECT object-name PROXY OUTGOING
```

- b. To enable OUTGOING PROXY access for the executor, use the following NCP command:

```
NCP> SET EXECUTOR OUTGOING PROXY ENABLED
```

2. Do the following on the target (remote) node:

- a. Run NCP, and use the following command to display the current settings for proxy access to the object.

```
NCP> SHOW OBJECT object-name CHARACTERISTICS
```

If PROXY INCOMING is not specified, and if proxy access is required for the requested object, use the following NCP command to enable incoming proxy access for this object only:

```
NCP> SET OBJECT object-name PROXY INCOMING
```

- b. Use the following command to display the current settings for proxy access to the remote executor:

```
NCP> SHOW EXECUTOR CHARACTERISTICS
```

If the executor INCOMING PROXY is not specified for the executor, then incoming proxy access to the executor is denied.

- c. To enable incoming proxy access to the executor, use the following NCP command:

```
NCP> SET EXECUTOR INCOMING PROXY ENABLED
```

Caution

Enabling proxy incoming on the executor can cause security problems. For more secure proxy access, set proxy incoming enabled only on the object.

- d. Run AUTHORIZE on the remote node, and use the following command to make sure that the proxy definitions in NETPROXY.DAT are correct:

```
UAF> SHOW/PROXY node::user
```

If the proxy definition for the source node is not correct, use the following command to change it:

```
UAF> MODIFY/PROXY node::user user
```

In this example, *node::user* is the source node and user, and *user* is the user on the target (or current) node.

The source node is the node that originates the proxy login request. The target node is the node that receives the proxy login request, in this case, the current node.

- e. Use the following AUTHORIZE command on the remote node to make sure that the proxy account the user is trying to access on the remote exists:

```
UAF> SHOW user-id
```

If the account does not exist, but should exist, use AUTHORIZE to create the account.

C. Do the following to resolve problems due to default access control information:

1. Log in to an account on the remote node that has SYSNAM and SYSPRV privileges, and perform the following steps.
2. Run NCP, and use the following command to check whether the object requested exists, and whether it has a user name and password associated with it:

```
NCP> SHOW KNOWN OBJECTS
```

3. If the object does not exist, create the object.
4. If the object exists, go to the next step.
5. Run AUTHORIZE, and use the following command to check that the object's user ID specified in the SHOW KNOWN OBJECTS command has an account in the remote node's system user authorization file (SYSUAF):

```
UAF> SHOW user-id
```

Login Information Invalid

6. If there is an account for the object's user ID, use the following AUTHORIZE command to define a new password for it:

```
UAF> MODIFY user-id/PASSWORD=password
```

7. Use the following NCP command to display the nonprivileged user ID and nonprivileged password for the remote node:

```
NCP> SHOW EXECUTOR CHARACTERISTICS
```

8. Run AUTHORIZE, and use the following command to check that the nonprivileged user specified in the executor characteristics has an account in the remote node's system user authorization file (SYSUAF):

```
UAF> SHOW user-id
```

9. If there is an account for the nonprivileged user, use the following AUTHORIZE command to define a new password for it:

```
UAF> MODIFY user-id/PASSWORD=password
```

10. If an account does not exist for the nonprivileged user, run AUTHORIZE, and use the ADD command to define an account.

11. Run NCP, and use the following command to define a nonprivileged user, and to specify the same password for the nonprivileged account that you defined in the SYSUAF:

```
NCP> SET EXECUTOR NONPRIVILEGED USER user-id PASSWORD password
```

Recommendations

You might want to use security alarms to provide information for troubleshooting login failure problems. Security alarms can provide information such as the user name and password used in failed login attempts. See the SET AUDIT command in the *VMS DCL Dictionary* for more information on setting up security alarms. Use the DCL command, REPLY/ENABLE=SECURITY, to display security alarms. Use of the REPLY/ENABLE=SECURITY command requires SECURITY privilege.

Because the security auditing features involve some system overhead, be careful to select the security features that provides the most benefit in your work environment. Overuse of alarm messages diminishes their usefulness; and, because alarm messages have priority over any other I/O, they can tie up the security operator's terminal.

Network Is Unreachable

Symptoms

Users on a TCP/IP network receive the following message when trying to connect to a host on a different network:

```
network is unreachable
```

Explanation

A host or IP router is sending the local host an ICMP message indicating that no path exists to the remote host's network. You can use the information from the ICMP message to help you understand how far your connection request traveled before it failed.

The problem is either on the local host or the path between the local and remote hosts. If the problem is on the local host, it may involve the local host's hardware, connection to the network, or routing tables. If the problem is not on the local host, it involves the path between the local and remote hosts.

Note

This message may not indicate a problem. Routers along the path to the remote host might have security features enabled that prevent you from reaching the remote host.

Troubleshooting Strategy

Make sure the following are correct:

- A. Configuration of the network devices on the local host
- B. Routing tables on the local host

Trace the path looking at each IP router's routing tables to ensure that there is an entry for the remote host's network. Repair the incorrect IP router's routing tables.

Note

This step requires a thorough knowledge of your topology.

- C. Local host's address-to-name translation for the remote host is correct.

Troubleshooting Procedure

- A. Make sure that the network devices are configured properly on the local host, using the following steps:

Network Is Unreachable

To check the configuration, you need to know the netmask and broadcast address for your network. The `/etc/ifconfig` command sets up the network devices. At system startup, the `/etc/rc.local` file configures the network devices.

1. Use the following command to display the configured network devices:

```
# netstat -i
```

2. If the necessary network device is not configured, configure it using either the `/etc/ifconfig` command or the `netsetup` command as follows:

- To configure the network device with the `/etc/ifconfig` command, use the following example as a guideline:

```
# /etc/ifconfig qe0 `bin/hostname` broadcast 16.0.255.255
netmask 255.255.0.0
```

This example configures a DEQNA for network 16, with the second octet of the address set for subnet addressing.

- To configure the network device with the `netsetup` command, log in to the superuser account.

For first time configurations, use the following command:

```
#/etc/netsetup install
```

For all existing configurations, use the following command:

```
#/etc/netsetup
```

The `netsetup` program prompts you for information about the remote host, and adds the information you supply to the `/etc/rc.local` file. The changes you make take effect when you reboot the system.

- B. Check the local host's routing tables, remembering that routing can occur through a host-specific route, a route specified for the destination network, or a default route.

See Figure 1-28 for a flow chart illustrating IP routing.

1. Use the following command on the local host to display the contents of the routing tables.

```
# netstat -r
```

If the routing tables show this routing information	Go to this step
A host-specific or destination network route	Step 2
A default route	Step 3
No route information	Step 4

2. If the routing tables show a host-specific or destination network route for the destination host, use the `ping` command to see if the IP router specified is reachable.

```
# ping IP_router_name
```

- If you cannot reach the IP router, make sure that the local host's cabling to the network is intact, and do the same for the IP router's cabling to the network.

Network Is Unreachable

- If you can reach the IP router, obtain the routing information from the router, and go to the table in step B1. The table in step B1 specifies what to do based on the type of routing information in the routing tables.
3. If the netstat command shows a default route, use the ping command to see if the default IP router is reachable:

```
# ping IP_router_name
```

- If the IP router is not reachable, make sure that the local host's cabling to the network is intact, and do the same for the IP router's cabling to the network.
 - If the IP router is reachable, obtain the routing table from the router, and go to the table in step B1. The table in step B1 specifies what to do based on the type of routing information in the routing tables.
4. If no route exists to the destination, add a route.

You can run the routing daemon to add routes automatically or use the route command to add the specific route manually to a router.

Using the routing daemon to add routes automatically

- a. Use the following command to see if the routing daemon (/etc/routed) is running:

```
# ps -aux | grep routed
```

The following example of output from this command shows that routed is running in quiet mode. ❶

```
root      77  0.0  0.8 204  88 ?  S  ❶ 4:42 /etc/routed -q
root     7255  0.0  0.3  40  32 p1 S    0:00 grep routed
```

- b. If the routing daemon is running, but there are still no routes, the local host is not receiving the routing updates.
Check the local host's cabling to the network.
- c. If the routing daemon is not running, but should be, run the routing daemon in quiet mode as follows:

```
# /etc/routed -q
```

- d. Make sure that the /etc/routed -q command is in the /etc/rc.local file.
- e. Wait a couple of minutes to allow for the routing tables to be filled and try to reach the remote host again.

If you still get the host is unreachable message, go to step B1, and repeat the procedure, now using the updated routing tables.

Network Is Unreachable

Using the route command to add a route manually

- a. To add a default route to a stable IP router, use the following command:

Note

Use an IP router that is only one hop away.

```
# route add default "ip_router_name" 1
```

- b. Try to reach the remote host again.
 - c. If you still get the host is unreachable message, go to step B1, and repeat the procedure.
- C. Make sure the local host's address-to-name translation for the remote host is correct.
1. Use the procedures shown in the "Unknown Host" problem in this chapter to check the address-to-name translation.
 2. Try to connect to the remote host.

Network Object Unknown

Symptoms

Users receive the following message when trying to access a remote node:

```
%SYSTEM-F-NOSUCHOBJ, Network object is unknown at remote node
```

Explanation

This is a DECnet-VAX node problem involving the session layer, and can be caused by any of the following:

- A. The object requested is not defined in NCP, is not started, or the file specified in the requested object does not exist.
- B. The user tried to access an object with an alias name, and ALIAS INCOMING is disabled for the object.
- C. If the user operation resulting in this message was SET HOST, SYSSYSTEM RTTLOAD.COM has not been run. As a result, the REMACP process is not running on the remote system, the drivers are not loaded, and the CTERM and REMACP objects do not exist.

Note

If the operation was SET HOST, and the REMACP process is not running, it may be that the system has just rebooted and has not yet executed RTTLOAD.COM. In this case, wait a few minutes, and try the operation again before beginning the troubleshooting procedure.

Troubleshooting Strategy

- A. For problems related to the requested object, do the following:
 1. Check to see if the object is defined in NCP.
 2. If the object is not defined, define it.
 3. Check to see if the file specified for the requested object exists.
 4. If the file for the requested object does not exist, create it.
 5. Check to see if the object is started.
 6. If the object is not started, start it.
- B. For problems related to the setting of ALIAS INCOMING on the requested object, do the following:
 1. Check the current setting of ALIAS INCOMING on the object.
 2. Enable ALIAS INCOMING on the object if necessary.
- C. For problems related to SET HOST and RTTLOAD.COM not running, do the following:
 1. Check if the REMACP process is running.

Network Object Unknown

2. Execute RTTLOAD.COM to run the REMACP process.

Troubleshooting Procedure

A. This step explains how to solve problems related to the requested object.

Objects can be associated with command and executable files, or they can be associated with a process. For objects associated with command and executable files, start with step 1. For objects associated with a process, start with step 5.

To determine the type of association the requested object has, run NCP and use the following command to see if the object is defined on the remote node:

```
NCP> TELL remote-node-id SHOW KNOWN OBJECTS
```

1. If the object is not defined, run NCP on the remote node, and define it using the following command and additional parameters, as required:

```
NCP> SET OBJECT object-id
```

2. If the object is defined, run NCP and use the following command to see if the object has a file specified:

```
NCP> TELL remote-node-id SHOW KNOWN OBJECTS
```

3. Use the DIRECTORY command to see if the file specified for the requested object exists.
4. If the file for the requested object does not exist, create it.
5. Use the following command on the remote node to see if the object is started:

```
$ SHOW SYSTEM
```

If the object's process is not displayed, the object is not started.

6. If the object is not started, start it.

The method for starting the object depends on which object needs to be started. For example, to start REMACP, execute the SYSSMANAGER:RTTLOAD.COM command procedure.

B. Do the following on each node in the cluster to resolve problems related to ALIAS INCOMING:

Note

Because the problem may be due to the improper setup of any of the cluster nodes, it is important to check the object on each node. If you do not resolve setup problems on each node, the problem may continue to occur intermittently, when an access request from a remote node reaches the improperly set up node.

1. Run NCP and use the following command to display the current setting for ALIAS INCOMING on the object:

```
NCP> SHOW OBJECT object-name CHARACTERISTICS
```

2. If ALIAS INCOMING is disabled for the object, enable it using the following command:

```
NCP> SET OBJECT object-name ALIAS INCOMING ENABLED
```

C. Do the following to address the problem if the user action was SET HOST and RTTLOAD.COM was not run:

1. Use the following command to see if the STARTNET job is running on the local node:

```
$ SHOW QUEUE batch-queue-id
```

The remote node may currently be unknown because the local node has not completed the STARTNET job. If the STARTNET job is still running, wait until it completes, then try the operation again. If the remote node is still unreachable, continue with the following steps.

2. Run NCP, and use the following command to determine if the REMACP process is running on the remote node:

```
NCP> TELL remote-node-id SHOW KNOWN OBJECTS
```

The NCP utility displays the list of objects running and their process identification numbers (PIDs). If the display does not include a PID for the REMACP object, the REMACP process is not running.

3. If the REMACP process is not running, execute the SYSSMANAGER:RTTLOAD.COM command procedure on the remote node.

The SYSSMANAGER:RTTLOAD.COM command procedure loads RTTDRIVER and CTDRIVER, and runs the REMACP process.

Network Partner Exited

Network Partner Exited

Symptoms

Users receive the following message when attempting to perform any network operation except SET HOST:

```
%SYSTEM-F-LINKEXIT, network partner exited
```

Explanation

This is a DECnet-VAX node problem involving the session layer. The error message indicates a problem on the remote node, potentially caused by any of the following:

- Improper protection on any of the following:
 - Requested object's executable and command files on the remote node
 - SYSSSYSTEM directory
 - NETSERVER.COM and NETSERVER.EXE files in the SYSSSYSTEM directory
 - Files pointed to by the SYSSSYLOGIN command file
 - Default DECnet directory
 - DCL tables in SYSSLIBRARY
- An error in the SYSSLOGIN command file
- An error in the default DECnet account's LOGIN.COM file (errors in this file force a logout)
- No LOGIN.COM file specified in the SYSUAF record for the default DECnet account (assuming the default DECnet account is captive)
Other parameters in the SYSUAF record for the default DECnet account may also cause this symptom, such as an insufficient value for the BYTLM parameter.
- A user attempting to access a disabled account on the remote node
- An error in starting the requested object
- The LOGIN.COM file for the account may specify some kind of interactive use
For example, the LOGIN.COM file may start a menu on login. This kind of interactive use may not work for remote DECnet connections.

Troubleshooting Strategy

1. Obtain the following information:
 - Operation the user was performing
 - Account the user attempted to access
 - Object the user was accessing
2. Check the following:
 - Setup of the account the user was accessing

- Protection on the object's executable and command files
3. Examine the NETSERVER.LOG file associated with the account for information explaining why the network link was aborted. Using the information in the NETSERVER.LOG file, you can correct the problem that caused the remote node to abort the link.

See VMS Netserver Log File for more information on the NETSERVER.LOG file.

Troubleshooting Procedure

1. Find out what operation the user was performing when the error occurred, and what account the user was accessing during the operation. The account might have been the user's account or the default nonprivileged DECnet account.

The type of access control the user specified helps you determine what account the user was accessing, for example:

- If the user specifies explicit access control information, access is through the account specified.
 - If the user did not specify explicit access control information, access is through the default DECnet account or through a proxy account.
 - When access is through the default DECnet account, the system uses the account associated with the object requested. If no account is associated with the object requested, access is through the nonprivileged user account specified for the executor.
 - If a proxy exists, access is through the account pointed to by the proxy.
2. Make sure that the accounts are set up properly on the remote node.

For example, for proxy accounts, do the following:

- a. Run AUTHORIZE, and use the following command to see if a proxy account is defined for the user's node and user name:

```
UAF> SHOW/PROXY local-node-id::user-id
```

- b. If a proxy exists for the account, run NCP and use the following commands to see if incoming proxy access is enabled for both the object and the executor. (See Table 5–6 for information on proxy access settings.)

```
NCP> SHOW OBJECT object-name CHARACTERISTICS  
NCP> SHOW EXECUTOR CHARACTERISTICS
```

NCP displays the current setting for incoming proxy access. If incoming proxy access is not enabled, then the system does not use the proxy to determine access. Instead, it grants access through the account associated with the object or the nonprivileged user account specified for the executor.

3. Log in to the remote node and use the following command to check the network objects on the remote node:

```
NCP> SHOW KNOWN OBJECTS CHARACTERISTICS
```

Note the names of the files specified for the requested object.

Network Partner Exited

4. Use the following command to check whether the executable and command files for the requested object on the remote node have the proper file protection:

```
$ DIRECTORY filename.exe, filename.com/PROTECTION
```

The file protection for the object's executable and command files should be world:read,execute. For example, if the user was attempting to copy a file, make sure that FAL.EXE and FAL.COM exist and have world:read,execute access specified.

5. If the file protection for world access is other than read, correct it using the following DCL command:

```
$ SET PROTECTION filename.exe, filename.com/PROTECTION=(W:R)
```

6. Use the following command to display the NETSERVER.LOG file for the account on this node:

```
$ DIRECTORY/DATE device:[directory]NETSERVER.LOG
```

If no NETSERVER.LOG file exists, go to Step 9.

7. Make sure that you have the correct NETSERVER.LOG file for the user's operation, and examine the NETSERVER.LOG file for errors.
Often, the problem may be with a login command file. Some common error messages are:

- "Error opening captive command procedure"
- "Duplicate process name"
- "Insufficient privilege or file protection violation"

8. Correct the error and tell the user to try the operation again.

9. If no current NETSERVER.LOG file exists, do the following:

- a. Make sure that you check the correct directory for the NETSERVER.LOG file.

- b. If you have checked the correct directory, but no current NETSERVER.LOG file exists, check directory protections and file protections to see if they are preventing a NETSERVER.LOG file from being created. Possible reasons that a NETSERVER.LOG file has not been created include the following:

- Improper protection on any of the following:
 - SYSSSYSTEM directory
 - Executable and command files in the SYSSSYSTEM directory for the object requested
 - NETSERVER.COM and NETSERVER.EXE files in the SYSSSYSTEM directory
 - Files pointed to by the SYSSSYSLOGIN command file
 - Default DECnet directory
 - DCL tables in SYSSLIBRARY
- The user attempting to access a disabled account
- No write access to the directory

Network Partner Exited

- No disk space
 - The existence of a NETSERVER.LOG file, version 32767
(NETSERVER.LOG;32767)
- c. Correct the problem preventing creation of the NETSERVER.LOG file.
 - d. Recreate the original user action.

If the original user action results in the "Network Partner Exited" message and the creation of a NETSERVER.LOG file, return to Step 7.

Node Out of Range Packet Loss

Node Out of Range Packet Loss

Symptoms

OPCOM displays one of the following messages:

Node out of range packet loss

Partial routing update loss

Explanation

These messages indicate a DECnet-VAX node problem, and usually occur when a new node joins the network. The messages result from one or both of the following:

- A. A node on the network has a DECnet address that is too high for the network area of which it is a part.

When this is the cause of the problem, the OPCOM messages usually occur on multiple nodes in the network.

A DECnet network can consist of a maximum of 63 areas, and each DECnet area can consist of a maximum of 1,024 nodes. However, a DECnet network can define a maximum number of areas that is less than 63, and each DECnet area can define a maximum number of nodes that is less than 1,024.

For example, assume a network has a maximum of 5 areas, with a maximum of 500 nodes per each area. In this case, all DECnet addresses in the area must be 500 or less. If the DECnet area number is 5, the range of node addresses for that area is 5.1 through 5.500.

If a node uses an area number that is greater than 5 or a node address greater than 500, the "node out of range packet loss" and "partial routing update loss" messages occur on nodes throughout the network when the incorrectly configured node announces its adjacency.

- B. The routing tables on one or more nodes in the network are not large enough to perform routing or end node updates for another node's valid DECnet address.

When this is the cause of the problem, the OPCOM messages usually occur on fewer nodes in the network than if the cause is an invalid (too high) DECnet address. In this case, the OPCOM messages occur only on the node or nodes that have insufficient routing tables.

Troubleshooting Strategy

Determine if the problem affects one or multiple nodes.

- A. If the problem affects many nodes on the network, check to see if a new node on the network has misdefined its DECnet address, and redefine the new node's address, if necessary.
- B. If the problem affects only one or a few nodes, check the routing tables on the affected node or nodes to see if the routing tables are adequate. Increase the size of the routing tables, if necessary.

Troubleshooting Procedure

A. If the problem affects many nodes, use the following steps to determine whether a node's DECnet address is too high, and to correct the address, if necessary.

1. Log in to any any node that displayed the OPCOM messages, and use the following command to enable the display of OPCOM messages on your terminal:

```
$ REPLY/ENABLE=NETWORK
```

OPCOM periodically displays various messages on your terminal, including the "node out of range" or "partial routing update packet loss" messages.

2. Run NCP and use the following command to display the local node's maximum address and maximum area definitions:

```
NCP> SHOW EXECUTOR CHARACTERISTICS
```

3. Make a note of the maximum address and maximum area definitions. Step 5 uses this information to help you identify the problem node.

4. Set the maximum address and maximum area to the highest values possible, as follows:

```
NCP> SET EXECUTOR MAXIMUM ADDRESS 1023  
NCP> SET EXECUTOR MAXIMUM AREA 63
```

Setting the maximum address and maximum area parameters as shown prevents any further "node out of range" or "partial routing update packet loss" messages.

Any nodes with DECnet addresses that were previously too high for the network now generate adjacency messages, which OPCOM displays on your terminal.

5. Look for an adjacency message for a node whose address is higher than your original maximum address and maximum area definitions. These messages indicate the node whose address is misdefined.
6. Get a correct address for the misdefined node from the person who assigns DECnet addresses on your network.
7. Log in to the misdefined node.
8. Run NCP and use the following commands to redefine the DECnet address correctly:

```
NCP> DEFINE EXECUTOR ADDRESS address  
NCP> SET EXECUTOR STATE OFF  
NCP> EXIT
```

9. Restart the network on the redefined node using the following command:

```
$ @SYS$STARTUP:STARTNET
```

10. Log in to the original node.

Node Out of Range Packet Loss

11. Use the following command to restart the network, thereby resetting the node's maximum address and maximum area parameters to their original settings:

```
$ @SYS$STARTUP:STARTNET
```

B. If the OPCOM message occurs on only one or a few nodes, do the following:

1. Log in to a node that is displaying the OPCOM message.
2. Run NCP and use the following command to display the local node's maximum address and maximum area definitions:

```
NCP> SHOW EXECUTOR CHARACTERISTICS
```

3. Use the following commands to increase the size of the routing tables on the local node:

```
NCP> DEFINE EXECUTOR MAXIMUM ADDRESS address  
NCP> SET EXECUTOR MAXIMUM ADDRESS address  
NCP> DEFINE EXECUTOR MAXIMUM AREA area  
NCP> SET EXECUTOR MAXIMUM AREA area
```

Partial Routing Update Loss

See "Node Out of Range Packet Loss"

Partitioned Area

Partitioned Area

Symptoms

A group of nodes in one area are unreachable from another area, or a significant number of areas are unreachable from an area. Users receive a variety of messages indicating that nodes or areas are not reachable, including "Path to network node lost" and "Remote node is not currently reachable," as well as various timeout messages.

With this symptom, some nodes in the unreachable area or areas may still be able to communicate with nodes in other areas.

Note

These symptoms do not necessarily indicate a partitioned area. What distinguishes the partitioned area problem from other related problems is that with partitioned areas, the group of nodes that is unreachable from a given location or locations can still be reached by other nodes on the network.

Explanation

This symptom indicates a DECnet WAN problem, which may be caused by one of the following:

- Failure of two or more area routers or circuits in an area, causing sections of an area to be unreachable.
- Failure to provide redundant circuits in an area. Redundant circuits help to ensure that no single point of failure exists between the level 2 routers in an area, by providing an alternate circuit to use in case one circuit fails. The path between level 2 routers cannot contain any level 1 routers.
- Failure of an area router's Ethernet controller.

Figure 5-3 shows how an area with redundant paths can become partitioned if two or more circuits fail. The dashed lines indicate the circuits that have failed. Because both of these circuits are unavailable, nodes V, W, and X are partitioned from nodes T, U, and Y.

Figure 5-3 Area Partitioned Due to Multiple Failures

Partitioned Area

Figure 5–4 shows how an area that lacks redundant paths can become partitioned if one or more circuits fail.

Figure 5–4 Area Partitioned Due to Configuration Weaknesses

In this figure, all circuit costs are equal to 1. The only path in Area 4 between the level 2 routers is through a level 1 router and circuits Y and Z. If all circuits are working, no problem exists. For example, node C in Area 3 attempts to communicate with node D in Area 4. If either circuit W or X fails, no problem arises because the remaining path into Area 4 provides a route to node D.

However, if circuits Y or Z fail, the level 2 router in Area 3 finds the path to the level 2 router in Area 4 on the basis of the least-cost algorithm: the path is from node C to node B to node A. Because circuit Y or Z is down, however, it is not possible to get to destination node D.

Another type of partitioned area problem occurs when a node in Area 4 tries to communicate with a node in Area 3, and circuit X is unavailable. The traffic cannot leave Area 4 because, with circuit X unavailable, there is no place for the traffic to go after it reaches the level 2 router, D. Circuit X is unavailable, and the traffic cannot revert to a level 1 router at this point. However, traffic coming into the area will be routed to the node without a problem. This is because incoming traffic is not subject to the same restrictions going through level 2 to level 1 routers as is outbound traffic.

Troubleshooting Strategy

The most important information to have in mind if you suspect a partitioned area problem on your network is what the network topology looks like. Specifically, you need to know the sites that make up the DECnet areas, the nodes used at each site for wide area routing, and the wide area circuits from the routers in an area to other sites and areas on the network.

You can use the Network Control Program (NCP) to query remote routers regarding the reachability of the suspected nodes and area. To query the correct routers, you must be familiar with the overall topology of your network.

To solve this problem, try to determine which areas can communicate with the suspected partitioned area, and which nodes within that partitioned area are reachable. By doing this, you isolate the cause of the failure and the partitioning, and you can focus on repairing the circuit or router that is causing the problem.

NMCC/DECnet Monitor can be especially helpful in isolating the cause of the problem.

Troubleshooting Procedure

1. Determine which area is unreachable.
2. Trace the routing path from various areas in the network to determine which routers and circuits or which circuits are down.
Look for circuits that are in the on-starting state, and for malfunctioning routers.
3. If it is possible to implement an interim solution that enables part of the partitioned area to communicate with the rest of the network, do so.
4. Repair the faulty circuits and routers, or faulty circuits.
5. Remove the interim solution.

Recommendations

Try to avoid configurations, such as the one illustrated in Figure 5-4, that use straight-line configurations (as in Area 4). Also, provide a direct link between all level 2 routers. For example, in Figure 5-4, installing a link between nodes A and D provides an alternative path for nodes in that area.

Permission Denied

Permission Denied

Symptoms

When using commands which are executed remote hosts (such as rsh and rcp), the user receives the following message, and the operation fails:

```
Permission denied
```

Explanation

This problem affects ULTRIX hosts and involves the internet protocol (IP).

ULTRIX systems determine whether to permit access for remote users through one of the following files:

- .rhosts file
- /etc/hosts.equiv file

When the permission denied message occurs, the problem may be due to one or more of the following:

- Incorrect host and user definitions in the user's .rhosts file on the remote host
- Improper setup of the /etc/hosts.equiv file
- Improper directory or file protection on files to be copied or the .rhosts file

Note

This symptom may not indicate a problem. It is possible that the remote host may be intentionally preventing remote access. Before you try to resolve this problem, be sure that the user is intended to have access to the remote host.

Troubleshooting Strategy

Make sure that the following conditions are met:

- The .rhosts file on the remote host contains the proper host and user definitions.
- The /etc/hosts.equiv file is set up properly.
- The directory and file protections are correct on the following files:
 - file to be copied
 - remote .rhosts file

Troubleshooting Procedure

Do the following on the remote host:

1. Display the contents of the `/etc/hosts.equiv` file, to determine if the user's host name is in that file:

```
# grep hostname /etc/hosts.equiv
```

- If the command returns you to the prompt, no entry exists in the `/etc/hosts.equiv` file for the host name you specified. If your site's security policy permits, you can edit the `/etc/hosts.equiv` file and add the host name.

Note

If you mistype the host name in the `grep` command, `grep` will not locate the host name even if the host name is in the file. Before you make any changes to the `/etc/hosts.equiv` file confirm that the host is not already in the `/etc/hosts.equiv` file.

- If the command displays an entry, make sure the host name is correct for the user.
 - If the entry is incorrect, modify the file to contain the correct definitions.
2. Use the following command to move to the user's login directory so you can check the `.rhosts` file:

```
# cd users_login_directory_name
```

3. Use the `ls` command to determine if a `.rhosts` file exists.
4. If the user's login directory does not contain a `.rhosts` file, use a text editor to create one that contains the correct user name and host name for the user.
5. If the user's login directory has a `.rhosts` file, use the following command to display its contents:

```
# grep hostname .rhosts
```

If the user or host name is missing or incorrect, modify the `.rhosts` file so that it contains the correct definitions.

Note

Be aware that some name services require the full domain name form of a host name, some require only a shortened form of the domain name. Be sure to use the proper form for the name service you use.

Also, if you are *not* running the domain name system with long names, your `/etc/hosts` file must not define host names with a long name directly following the IP address. If you are using short host names, then make sure your `/etc/hosts` file has short names only.

Permission Denied

6. Make sure that the local host knows about the remote host.
Although the local host may receive the message, permission denied, the true cause of the problem may be that the remote host is unknown. See "Unknown Host" in this chapter for more information on how to correct this problem.
7. Use the following command to confirm that you are in the user's login directory:

```
# pwd
```
8. From the user's directory, use the following command to check the directory and file protections on any files the users wants to access (including .rhosts):

```
# ls -l file-name
```

The user needs read access at the file level.
9. If the files do not have read privilege, use the following command to change the file and directory protection:

```
# chmod u+r file-name
```
10. Use the following command to display the protection on the directory:

```
# ls -l -d
```

The user needs read and write privilege to use rcp.
11. If the directory does not have read and write privilege, use the following command to change the file and directory protection:

```
# chmod u+w,r directory-name
```

Remote Node Is Not Currently Reachable

Symptoms

A user gets the following message when attempting any network operation:

```
%SYSTEM-F-UNREACHABLE, remote node is not currently reachable
```

This message occurs when a user tries to establish a connection to a remote node, or when a user tries to reestablish a connection that was lost. This message may also occur when a user tries to connect to a cluster system using the cluster alias. The connection fails when the user specifies the cluster alias, but succeeds if the user specifies a particular node in the cluster.

Explanation

This symptom indicates a node, LAN, or WAN problem, and means that the remote node is not reachable through DECnet. This symptom may be due to any of the following:

- The remote node is not running.
- The lines and circuits may not be operating properly on the local or remote node.
- A routing node may not be running or operating properly.
- A problem exists on the routing path.
- The remote node may be incorrectly defined in the local node's NCP node database.
- The local node may have setup problems such as the following:
 - Circuit, line, or both may not be defined.
 - Values for MAXIMUM HOPS, MAXIMUM COSTS, and MAXIMUM VISITS might be too low.
- If the remote system is a cluster system, there may be setup problems such as the following:
 - MAXIMUM BROADCAST NONROUTER parameter for the routing node is set too low.
 - Cluster system has not defined one of the nodes to be the routing node for the cluster. In a cluster system, at least one node has to be the routing node.
 - ENQUEUE limit on NETACP on the routing system is too low.

Troubleshooting Strategy

The troubleshooting procedure explains how to solve the problem if it is on the local node, the routing path, or the remote node.

- A. Check the following on the local node:
- Remote node's address in the local node's volatile node database
 - Local node's circuits and lines

Remote Node Is Not Currently Reachable

- Values of MAXIMUM HOPS, MAXIMUM COSTS, and MAXIMUM VISITS
 - ENQUE limit values for NETACP
- B. Trace the routing path to determine if the problem is on the routing path or the remote node.
- If the problem is on the routing path, do loopback tests or reachability tests (as described in NCP Loopback Tests and Reachability Tests for TCP/IP Networks) to determine the cause.
- C. If the problem is on the remote node, check the following on the remote node:
- Remote node is running.
 - DECnet is running.
 - Circuits and lines are operating properly.
 - Cluster aliasing is set up properly.
 - ENQUE limit for NETACP is correct.

Troubleshooting Procedure

- A. **Do the following to determine if the local node is the cause of the problem:**
1. Use the following steps to determine if the remote node's node address is correctly defined in the local node's volatile node database, and to redefine it, if necessary.
 - a. Determine if the node address for the remote node is properly defined by checking the local node's definition against the master list of node definitions maintained for your network. Usually a single node (called *master-list-node* in the following example) maintains the master list of node definitions.

Run NCP, and use the following commands to get this information:

```
NCP> SHOW NODE node-id
NCP> TELL master-list-node SHOW NODE node-id
```
 - b. If the preceding commands display different node definitions for the remote node, use the following command to define the node address to the definition displayed by the master-list-node:

```
NCP> CLEAR NODE node-id ALL
NCP> PURGE NODE node-id ALL
NCP> DEFINE NODE node-address NAME node-id
NCP> SET NODE node-id ALL
```

Remote Node Is Not Currently Reachable

2. Use the following steps to check the lines and circuits on the local node:
 - a. Run NCP and use the following commands to determine the status of the lines and circuits:

```
NCP> SHOW KNOWN LINES
NCP> SHOW KNOWN CIRCUITS
```

- b. If the circuit state is on-starting, go to the corrective action for "Circuit State Problems" in this chapter.
 - c. If the lines and circuits are not on, use the following commands to turn them on:

```
NCP> SET LINE line-id STATE on
NCP> SET CIRCUIT circuit-id STATE on
```

3. Check the MAXIMUM HOPS, MAXIMUM COSTS, and MAXIMUM VISITS values, and modify them, if necessary, using the following NCP commands:

- a. Use the following command to display the current values for the MAXIMUM HOPS, MAXIMUM COSTS, and MAXIMUM VISITS:

```
NCP> SHOW EXECUTOR CHARACTERISTICS
```

A value of 15 for MAXIMUM HOPS is usually sufficient, unless your network is particularly wide. The MAXIMUM VISITS value should be two times the value of MAXIMUM HOPS.

Note

Use care in increasing the MAXIMUM HOPS value because a higher MAXIMUM HOPS value on routing systems can cause increased routing traffic. The increased traffic can cause the network to take longer to stabilize when a system stops running.

- b. Use the following NCP commands to increase the MAXIMUM HOPS, MAXIMUM COSTS, and MAXIMUM VISITS values:

```
NCP> SET EXECUTOR MAXIMUM HOPS value
NCP> SET EXECUTOR MAXIMUM COSTS value
NCP> SET EXECUTOR MAXIMUM VISITS value
```

4. If the local node is part of a Local Area VAXcluster, the problem may be related to ENQUE limits. Do the following to resolve problems related to ENQUE limits:

- a. Use the following command to display the current ENQUE limit value:

```
$ SHOW LOGICAL NETACP$ENQUEUE_LIMIT
```

The ENQUE limit on the routing node for NETACP must be two times the number of satellites plus 10.

- b. If the ENQUE limit is too low (as determined in step a), and DECnet *has not* started, use the following command to increase the value, otherwise, go to step c:

```
$ DEFINE/SYSTEM NETACP$ENQUEUE_LIMIT xxx
```

Remote Node Is Not Currently Reachable

- c. If the ENQUE limit is too low (as determined in step a), and DECnet *has* started, use the following commands to increase the value:

```
$ DEFINE/SYSTEM NETACP$ENQUEUE_LIMIT xx
$ MCR NCP
NCP> SET EXECUTOR STATE OFF
NCP> EXIT
$ @STARTNET.COM
```

B. Trace the routing path to determine if the problem is on the routing path or on the remote node.

If the problem is on the routing path, use loopback, modem, and circuit tests on the routing path (at the disconnect point) to help determine the source of the problem. (See NCP Routing Path Trace Procedure for information about these tests.)

If the problem is on the remote node, go to the next step.

C. If the problem is on the remote node, do the following:

1. Make sure that the remote node is running.
2. On the remote node, use the following command to make sure that DECnet is running:

```
$ SHOW NETWORK
```

3. On the remote node, use the following steps to check the lines and circuits on the node:

- a. Run NCP, and use the following commands to make sure that the lines and circuits are on:

```
NCP> SHOW KNOWN LINES
NCP> SHOW KNOWN CIRCUITS
```

- b. If the circuit is on-starting, go to the corrective action for "Circuit State Problems" in this chapter.

- c. If the lines and circuits are not on, use the following command to turn them on:

```
NCP> SET LINE line-id STATE ON
NCP> SET CIRCUIT circuit-id STATE ON
```

4. If the remote node is part of a VAXcluster, use the following steps to resolve problems related to the MAXIMUM BROADCAST NONROUTER parameter:

- a. Run NCP, and use the following command to display the value for the MAXIMUM BROADCAST NONROUTER parameter:

```
NCP> SHOW EXECUTOR CHARACTERISTICS
```

The value for MAXIMUM BROADCAST NONROUTER specifies the number of nonrouting nodes (end nodes) the executor node can have on its Ethernet circuits. The value for MAXIMUM BROADCAST NONROUTER should be at least the number of nonrouting (end nodes) on the Ethernet. The default value is 64.

- b. If it is necessary to increase the MAXIMUM BROADCAST NONROUTER parameter, use the following NCP command:

```
NCP> SET EXECUTOR MAXIMUM BROADCAST NONROUTER n
```

Remote Node Is Not Currently Reachable

5. If the remote node is part of a VAXcluster, use the following commands to resolve problems resulting from improper setup of the cluster:

- a. Run NCP, and use the following command on each node in the cluster to display the routing status of each:

```
NCP> TELL NODE node-id SHOW EXECUTOR CHARACTERISTICS
```

Check the "type" displayed for the node. The cluster must have a router defined (designated as ROUTING IV) and the router must be running for the cluster alias to work.

- b. If the cluster does not have a router defined, define one, using the following commands:

- i Shut down the network on the cluster using the following NCP command:

```
NCP> SET EXECUTOR STATE OFF
```

- ii Change the executor type using the following command:

```
NCP> DEFINE EXECUTOR TYPE ROUTING IV
```

- iii Execute the STARTNET.COM file to restart the network:

```
$ @STARTNET.COM
```

- iv If the router is defined but not running, restart it or define another node to be the routing node.

- c. If the remote node is part of a Local Area VAXcluster, the problem may be related to ENQUE limits. Use the following commands to resolve problems related to ENQUE limits:

- i Check LOADNET.COM for the current value for ENQLM.

The ENQUE limit on the routing node in NETACP must be two times the number of satellites plus 10.

- ii If the ENQUE limit is too low, and DECnet *has not* started, use the following command to increase the value; otherwise, go to step c:

```
$ DEFINE/SYSTEM NETACP$ENQUEUE_LIMIT xx
```

- iii If the ENQUE limit is too low (as determined in step a), and DECnet *has* started, use the following commands to increase the value:

```
$ DEFINE/SYSTEM NETACP$ENQUEUE_LIMIT xx
$ MCR NCP
NCP> SET EXECUTOR STATE OFF
NCP> EXIT
$ @STARTNET.COM
```

Terminal Server Connection Failures

Terminal Server Connection Failures

Symptoms

A user cannot connect to a service, VMS node, or ULTRIX host from a terminal server.

Explanation

This symptom indicates a LAN problem involving LAT protocol. The terminal server cannot connect to the requested service, node or host because of one of the following reasons:

- The group code is undefined, or the group codes between the terminal server and the requested service, node or host do not match.
- The terminal server node limit may be exceeded.
- LAT protocol may not be started on the requested service, node, or host.
- Resources are insufficient on the requested service, node, or host.

Troubleshooting Strategy

To troubleshoot this problem, do the following:

- A. Make sure that the group code definitions on the user's terminal server port match those on the requested service, VMS or ULTRIX system.
- B. Check the server node limit defined on the terminal server, and increase it, if necessary.
- C. Make sure that the LAT protocol is started on the requested service, VMS or ULTRIX system.
- D. Make sure that sufficient resources exist on the VMS or ULTRIX system.

Note

You can correct this problem using TSM or NCP to enter commands on the terminal server, or you can go directly to the terminal server and enter the commands.

Whether you use TSM or NCP, or enter the commands directly on the terminal server, the procedure is the same. Only the prompts change, depending on the method you use. For the purposes of illustration, TSM> is the prompt for the server commands in the following steps.

Troubleshooting Procedure

A. Make sure that the group code definitions on the user's terminal server port match those on the requested service, VMS or ULTRIX system.

1. Log in to the terminal server, using one of the following methods:

- To use TSM to log in to the terminal server, run TSM and use the following command:

```
TSM> USE SERVER server-id
```

- To use NCP to log in to the terminal server, the terminal server must be defined in the NCP database. If the terminal server is defined in the NCP database, run NCP and use the following command to log in to it:

```
NCP> CONNECT NODE server-id
```

If the terminal server is not defined, you can use the following NCP command:

```
NCP> CONNECT VIA PHYSICAL ADDRESS ethernet-physical-address
```

When you enter the following terminal server commands, the system displays the Local> prompt rather than the TSM> prompt. To return to NCP from the Local> prompt, press Ctrl/D.

- To enter commands directly on the terminal server, go to the terminal server and log in using your user name.

```
Enter username> username
```

When you enter the following terminal server commands, the system displays the Local> prompt rather than the TSM> prompt.

Note

From this point on, the commands are the same regardless of how you access the terminal server. However, the following prompts assume you use TSM to access the terminal server.

2. For DECserver 100 terminal servers, use the following command to display the group codes defined on the terminal server:

```
TSM> SHOW SERVER
```

For DECserver 200 terminal servers, use the following commands to display the group codes defined on the terminal server:

```
TSM> SHOW SERVER server-id  
TSM> SHOW PORT port-id
```

Terminal Server Connection Failures

3. If the requested service is a VMS system, do the following to display the group codes definitions on that system:

- a. Run LATCP on the user's requested service, and use the following command to check the group codes:

```
LCP> SHOW CHARACTERISTICS
```

- b. Compare the group code definitions displayed in LATCP to those in the LTLOAD.COM file. The definitions should be the same.
- c. Compare the group code definitions on the VMS system to those you displayed in step 2 on the terminal server.
 - If the definitions are *not* the same for the VMS system and the terminal server, go to step 5.
 - If the definitions are the same for the VMS system and the terminal server, go to step B.

4. If the requested service is an ULTRIX system, do the following to display the group code definitions on that system:

- a. Use the following LAT control program (lcp) command to display the host characteristics:

```
# lcp -d
```

- b. Compare the group code definitions displayed in lcp to those in the rc.local file. The definitions should be the same.
- c. Compare the group code definitions on the ULTRIX system to those you displayed in step 2 on the terminal server.
 - If the definitions are *not* the same for the ULTRIX system and the terminal server, go to step 5.
 - If the definitions are the same for the ULTRIX system and the terminal server, go to step B.

5. Determine whether the group code definitions are correct on the terminal server or the service node (VMS or ULTRIX system), and correct as necessary.

Step a explains how to redefine the group code definitions on the terminal server.

Step b explains how to redefine the group code definitions on the service node (VMS or ULTRIX system).

- a. If the service node (VMS or ULTRIX system) definitions are correct, then redefine the terminal server definitions to match the service node definitions as follows:
 - i Use the following commands to redefine the group codes on the terminal server to match the group codes on the VMS or ULTRIX service node:

Note

These commands require that you use the SET PRIVILEGED command on the terminal server to enable privileges for these operations.

Terminal Server Connection Failures

```
TSM> SET SERVER/ENABLE=group-list
TSM> DEFINE SERVER/ENABLE=group-list
```

- ii Use the following commands to define the port characteristics:

```
TSM> SET PORT AUTHORIZE GROUP group-id
TSM> DEFINE PORT AUTHORIZE GROUP group-id
TSM> SET PORT GROUP group-id
TSM> DEFINE PORT GROUP group-id
```

- b. If the terminal server definitions are correct, then redefine the service node (VMS or ULTRIX) definitions to match the terminal server definitions, as follows:

For VMS systems:

- i Run LATCP, and use the following commands to redefine the group codes on the service node:

```
LCP> SHOW CHARACTERISTICS
LCP> SET NODE/ENABLE=GROUPLIST
LCP> SET NODE/DISABLE=GROUPLIST
```

- ii Edit the SYS\$MANAGER:LTLOAD.COM file to reflect the same group codes you just specified or deleted in LATCP.

For ULTRIX systems:

- i Use the following lcp commands to redefine the group codes on the service node, substituting the appropriate group numbers for group_n:

```
# lcp -d
# lcp -g group_n, group_nn, group_nnn...
```

- ii Edit the /etc/rc.local file to reflect the same group codes you just specified or deleted in lcp.

- B. **If the group codes are properly defined and the problem persists, the number of systems defined in the group codes may be greater than the server node limit on the terminal server. Do the following:**

1. Use the following command to display the server node limit on the terminal server:

```
TSM> SHOW SERVER
```

2. If the server node limit on the terminal server is too small to accommodate the number of systems the terminal server needs to access, increase it using the following commands:

```
TSM> SET SERVER NODE LIMIT n
TSM> DEFINE SERVER NODE LIMIT n
```

- C. **Make sure that the LAT protocol is started on the requested VMS or ULTRIX system.**

For VMS systems:

Use the following command to start the LAT protocol:

```
$ @LTLOAD.COM
```

Terminal Server Connection Failures

For ULTRIX systems:

1. Make sure that the configuration file contains the following lines:

```
options          LAT
pseudo-device   lat
pseudo-device   lta n
```

For non-RISC systems, the configuration file is `/sys/conf/"hostname"`. For RISC systems, the configuration file is `/sys/conf/mips/"hostname"`.

For "lta n," specify the number of terminals, using a multiple of 16.

2. Add the following lines to the `/etc/ttys` file to define the tty devices:

```
tty08  "/etc/getty std.9600"  vt100  on nomodem #LAT
.
.
.
tty24  "/etc/getty std.9600"  vt100  on nomodem #LAT
```

3. Use the following `lcp` command to make sure that the LAT protocol starts with the appropriate group code designations:

```
# lcp -s -G groupn, groupn...
```

4. Make sure that the `/etc/rc.local` file contains the following line so that the LAT protocol starts, and has the appropriate group codes defined:

```
lcp -s
```

5. Make sure that DECnet or IP is running before you try to start LAT on an ULTRIX system.

ULTRIX requires DECnet or IP to be running before LAT can start.

D. **Make sure that sufficient resources exist on the VMS or ULTRIX system.**

For VMS systems, see "Insufficient Resources at Remote Node" in this chapter.

For ULTRIX systems, be sure that enough ports exist. If all ports are in use, increase the number of ports by a multiple of 16.

Unknown Host

Symptoms

Users on TCP/IP networks receive one of the following messages when trying to access a remote host:

```
unknown host
```

Explanation

This is an ULTRIX host problem, where the remote host's name has not been translated to an IP address. The local host could be using any or all of the following means for name-to-address translation:

- /etc/hosts file
- BIND name service
- YP name service

Note

This symptom may not indicate a problem. It is possible that the user is trying to reach the remote host using a name that is not assigned. Also, if the remote host is in your name domain, you can specify just the remote host name. However, if the remote host is in another name domain, you must specify the full domain name.

Troubleshooting Strategy

To solve the "Unknown Host" problem, do one or both of the following:

- A. Add to or change the order of the name-to-address services listed in the /etc/svcorder or /etc/svc.conf files on the local host so that the name-to-address translation for the remote host occurs.
- B. Add the correct translation for the remote host to the service that provides name-to-address translation for the local host.

Troubleshooting Procedure

A. Modify the name-to-address services listed in the /etc/svcorder or /etc/svc.conf files on the local host so that the name-to-address translation for the remote host occurs.

1. Check the /etc/svcorder or /etc/svc.conf file to determine how the local host obtains name-to-address translation information:

The local host could use any or all of the following:

- /etc/hosts file (local)
- BIND name service (bind)
- Yellow Pages name service (yp)

Unknown Host

The `/etc/svcorder` file for ULTRIX versions up to Version 3.1 specifies the methods the local host uses for all name-to-address translations, and the order in which the local host uses the methods.

The `/etc/svc.conf` file for ULTRIX Version 4.0 specifies the methods the local host uses for name-to-address translations for each type of name translation request, and the order in which the local host uses the methods.

For example, the `/etc/svc.conf` file specifies how the system translates host names, alias names, group names, network names, and so forth. The "host" field lists the methods used to translate host names and the order in which the methods are used.

2. If your file lists local as the only name-to-address translation mechanism, then the cause of the problem is that the `/etc/hosts` file does not have information on the remote host. Go to step B.
3. If your site uses BIND or YP name services for name-to-address translation, check to see if the BIND or YP services have information about the remote host.

For a BIND name server, do the following:

- a. Use the `nslookup` command to query the BIND name servers for the IP address of the remote host, as shown:

```
# nslookup
> hostname
```

Note

If the remote host is in your name domain, you can specify just the remote host name. However, if the remote host is in another name domain, you must specify the full domain name, and end your query with a period, as follows:

```
> hostname.csdept.stateu.edu.
```

-
- b. If the `nslookup` command displays the full domain name and IP address for the remote host, do one of the following:
 - Update the `/etc/svcorder` or `/etc/svc.conf` file on the local host to include the BIND name server as an available service.
 - Use the information displayed to update the `/etc/hosts` file on the local host as described in step B.
 - c. If the `nslookup` command displays a message stating that the BIND name server does not have information about the remote host, do one of the following:
 - Check to see if the Yellow Pages (YP) name service is available for your site, and go to the instructions for Yellow Pages name servers.
 - If no other services are available, obtain the translation information from a user on the remote host or through the local registrar, and go to step B.

- d. If the `nslookup` command generates no response from the BIND name service, check the `/etc/resolv.conf` file on the local host to determine which BIND name servers the host uses, and in what order the host uses them.
- e. It is possible that the bind server is not available. If the `/etc/resolv.conf` file does not show a bind name server, modify the `/etc/resolv.conf` file by adding the appropriate bind server information.
- f. Repeat steps a-c.
- g. If the bind server is not available, obtain the translation information from a user on the remote host or through the local registrar, and go to step B.

For a Yellow Pages name server, do the following:

- a. Use the following command to display the name-to-address translation database files:

```
% ypcat -x
```

The `ypcat -x` command displays the database names for various types of name-to-address translations, such as hosts, networks, groups and aliases. To solve "Unknown Host," use the database file specified for host name translations. In the following example, the file is the "hosts" file.

- b. Use the following command to query the YP name server for information about the remote host:

```
% ypcat hosts | grep hostname
```

- c. If the `ypcat` command displays the remote host's name and IP address, do one of the following:
 - Update the `/etc/svcorder` or `/etc/svc.conf` file to include the YP name server as an available service.
 - Use the information displayed to update `/etc/hosts` file, as described in step B.

B. If the `/etc/hosts` file or the name servers do not contain translation information for a remote host, add the translation information to the `/etc/hosts` file or the name servers.

For the `/etc/hosts` file, do the following:

To add the host name and IP address, edit the `/etc/hosts` file or run `netsetup`, the network set up program.

To use `netsetup`, log in to the superuser account, and type the following command:

```
#/etc/netsetup
```

The `netsetup` program prompts you for information about the remote host, and adds the information you supply to the `/etc/hosts` file. The changes you make take effect immediately.

For BIND or YP name servers, do the following:

Ask the manager of these servers to update the databases to include the correct information for the remote host.

Unknown Host

Recommendations

If you find that the BIND or YP name servers available for your site do not have information for remote hosts that you need to reach, notify the BIND or YP server managers, so that they can decide whether to include the remote hosts' information in the server databases.

Verification Reject

Symptoms

A circuit alternates between the on-starting and on-synchronizing states, and OPCOM displays verification reject messages, such as the following. This problem is most common with dialin connections to DECnet networks.

```

                %%%%%%%%% OPCOM 19-OCT-1989 15:37:07.40 %%%%%%%%%
                Message from user DECNET on NODE1
                DECnet event 4.6, verification reject
                From node 56.689 (NODE1), 19-OCT-1989 15:37:07.35
                Circuit TX-0-6, Node = 56.1014 (NODE2)

                %%%%%%%%% OPCOM 19-OCT-1989 15:37:20.65 %%%%%%%%%
                Message from user DECNET on NODE1
                DECnet event 4.6, verification reject
                From node 56.689 (NODE1), 19-OCT-1989 15:37:20.59
                Circuit TX-0-6, Node = 56.1014 (NODE2)

                %%%%%%%%% OPCOM 19-OCT-1989 15:37:30.66 %%%%%%%%%
                Message from user DECNET on NODE1
                DECnet event 4.6, verification reject
                From node 56.689 (NODE1), 19-OCT-1989 15:37:30.59
                Circuit TX-0-6, Node = 56.1014 (NODE2)

```

Explanation

This symptom indicates a DECnet-VAX node problem. The DECnet routing layer provides a means for verifying passwords between adjacent nodes, known as *circuit verification*. Because circuit verification can disallow access from one node to another, it provides additional security for DECnet nodes.

You can set passwords and enable circuit verification on a node-by-node basis. Because of this, nodes can be adjacent to a node through different circuits, and multiple nodes can be adjacent to a node through a single circuit.

A node may receive a verification reject message from a node to which it is intended to have access, because the passwords between the two nodes have been improperly defined.

Troubleshooting Strategy

1. Determine whether the two nodes in question are intended to have verification enabled.
2. If verification is not required, disable it.
3. If verification is required, ensure that the transmit and receive passwords for each node are properly defined.

Verification Reject

Troubleshooting Procedure

1. Find out if the verification is required for the two nodes.
2. If verification is not necessary, run NCP, and disable verification on each node using the following command:

```
NCP> SET CIRCUIT circuit-id VERIFICATION DISABLED
```

3. If verification is required, make sure that the local and remote node passwords are properly defined.
 - a. On each node, use the following command to display the other node's executor characteristics, including transmit and receive passwords, if any exist.

```
NCP> SHOW NODE remote-node-id CHARACTERISTICS
```

If the remote node has transmit or receive passwords defined, the local node's transmit password must match the remote node's receive password. Likewise, the local node's receive password must match the remote node's transmit password.

- b. Use the following commands to set up coordinated transmit and receive passwords on the local and remote nodes.

On the first node:

```
NCP> SET NODE remote-node TRANSMIT PASSWORD password_a-  
_NCP> RECEIVE PASS password_b
```

On the second node:

```
NCP> SET NODE remote-node TRANSMIT PASSWORD password_b-  
_NCP> RECEIVE PASSWORD password_a
```

Ethernet Configuration Guidelines

This appendix provides configuration guidelines for the following:

- Baseband Ethernet
- ThinWire Ethernet
- Fiber-optic cable
- DELNI Local Network Interconnect
- ThinWire multipoint repeaters (DEMPRs)
- Ethernet repeaters (DEREPs)
- Bridges

Baseband Ethernet

Configuration guidelines for Baseband Ethernet are as follows:

- 500 meters maximum per length of coaxial cable segment.
- 500 meters maximum per length of coaxial cable segment when connected with a DELNI and a DEMPR.
- 300 meters maximum per length of coaxial cable segment if you cascade a DEMPR or DELNI connected on the Ethernet segment.
- 100 taps maximum per coaxial cable.
- 50 meters maximum per transceiver cable from the board to the coaxial cable.
- 1023 stations maximum per nonextended LAN.
- 1500 meters maximum of coaxial cable between any two stations in a nonextended LAN.
- 2.5 meters minimum separation between stations.
- Use DESTAs to adapt ThinWire to baseband Ethernet devices.
- Use DESPRs and DEMPRs to adapt baseband to ThinWire Ethernet devices.

ThinWire Ethernet

Configuration guidelines for ThinWire Ethernet are as follows:

- 185 meters per ThinWire segment.
- 30 connections maximum per ThinWire cable.
- 0.5 meters minimum separation between stations.
- No cabling is allowed between the T-connector and the DESTA.

Ethernet Configuration Guidelines

ThinWire Ethernet

- No cabling is allowed between the T-connector and the Ethernet controller.

Fiber-Optic Cable

Configuration guidelines for Fiber-optic cable are as follows:

- 1000 meters maximum of fiber-optic cable total within a nonextended LAN
- 1500 meters maximum of fiber-optic cable between a remote bridge and a remote repeater
- 10,000 meters maximum of fiber-optic cable between two remote bridges
- 1000 meters maximum of fiber-optic cable between two stations within a nonextended LAN

DELNI Local Network Interconnect

Configuration guidelines for the DELNI local network interconnect are as follows:

- DELNIs *cannot* be cascaded if connected to an H4000 transceiver in GLOBAL mode.
- DELNIs *can* be cascaded if the top-level DELNI is in LOCAL mode.

DEMPR ThinWire Multiport Repeaters

Configuration guidelines for DempR ThinWire Multiport Repeaters are as follows:

- DempRs provide grounding and termination on one end for ThinWire segments.
- DempRs support a maximum of 29 stations on each segment.
- DempRs must be connected to an H4000-ba (no heartbeat) transceiver, or a standalone DELNI in GLOBAL mode with a loopback connector in the global port.

DEREP Ethernet Repeaters

Configuration guidelines for DEREp Ethernet Repeaters are as follows:

- DEREps must be connected to an H4000-aa (with heartbeat) transceiver.
- Two repeaters (DEREP, DempR, DESPR) maximum are allowed between any two stations (this is the *two repeater rule*).
- A pair of remote repeaters connected to either end of a fiber-optic cable count as one repeater.
- A remote repeater and remote bridge combination counts as one repeater and one bridge.
- Parallel (or redundant) repeaters are allowed in LOCAL mode only, set using the standby switch on the repeater.

Bridges

Configuration guidelines for bridges are as follows:

- A maximum of seven bridges are allowed between any two stations.
- A maximum of 8000 stations are allowed on an extended LAN.
- Bridges effectively reset the two repeater rule. Therefore, if you use a bridge between sets of repeaters, you can have more than two repeaters between two stations.
- Parallel (or redundant) bridge configurations are permitted. The spanning tree algorithm automatically disables one bridge to prevent loops. Or, you can use RBMS to disable a bridge and prevent loops.
- A pair of remote bridges on either end of a fiber-optic cable count as two bridges.

RFC Request Process

Requests for Comments (RFCs) are a series of technical papers that include proposals and protocol standards for the Internet. This appendix explains how to get copies of these documents through the following:

- Internet File Transfer Protocol (ftp)
- Network Information Center Automatic Mail Service
- Network Information Center Telephone Request

Obtaining RFCs Through The Internet File Transfer Protocol

This section explains how to do the following:

- Log into a remote host from which you can obtain RFCs
- List the files contained in public directories
- Get an index of available RFCs

Logging In To The Remote Host

1. Log in to nic.ddn.mil using the Internet File Transfer Protocol (ftp), as follows:

```
% ftp nic.ddn.mil
```

The remote host displays a message such as the following:

```
Connected to nic.ddn.mil.  
220 NIC.DDN.MIL FTP Server Process 5Z(47)-6 at Fri 22-Jun-90 11:04-PDT
```

2. Enter "anonymous" for the username:

```
Name (nic.ddn.mil:myname): anonymous
```

3. Enter "guest" for the password:

```
Password (nic.ddn.mil:anonymous):guest
```

The remote host displays a message such as the following:

```
331 ANONYMOUS user ok, send real ident as password.  
230 User ANONYMOUS logged in at Fri 22-Jun-90 11:04-PDT, job 39.
```

RFC Request Process

Obtaining RFCs Through The Internet File Transfer Protocol

Using Public Directories on The Remote Host

After you log in to nic.dnn.mil, you can use ftp to display a list of the files available in public directories. Table B-1 lists the public directories and a brief description of their contents.

Table B-1 Public Directories on NIC.DDN.MIL

Public Directory	Contents
DDN-NEWS:	DDN management bulletins and DDN newsletters
IEN:	Internet experimental notes
INTERNET-DRAFTS	Internet Engineering Task Force (IETF) idea papers
NETINFO:	Administrative notes, documents, host/TAC information, domain and Internet information, PC/Kermit information, NSC/HA lists, network program information, DDN Vendors Guide, and DDN New Users Guide
NETPROG:	Network programs for WHOIS, HOSTNAME, getting the NIC host table and host table compiler
PROTOCOLS:	Network protocols information
RFC:	Requests for comments

Use the ftp directory command, dir, to list the file names in the public directories. For example, the following command lists the file names in the RFC directory:

```
ftp> dir RFC:
```

Obtaining an Index of Available RFCs

Each public directory also contains an index or indexes to the files in that directory. For the RFC directory, the file that contains the index is RFC-INDEX.TXT. The following example shows how to get a copy of this file:

```
ftp> get RFC:RFC-INDEX.TXT
200 Port 12.83 at host 130.180.4.1 accepted.
150 ASCII retrieve of TS:<RFC>RFC-INDEX.TXT (56 pages) started.
226 Transfer completed. 143171 (8) bytes transferred.
local: RFC:RFC-INDEX.TXT remote: RFC:RFC-INDEX.TXT
143171 bytes received in 47 seconds (3 Kbytes/s)
ftp>
```

RFC Request Process

Obtaining RFCs Through The Internet File Transfer Protocol

Copying RFC Files

To get a copy of an RFC, use the following command syntax:

```
ftp> get RFC:RFCnnnn.type
```

"nnnn" is the number of the RFC you want

"type" is either .PS for the postscript version, or .TXT for the ASCII text version

The following example shows how to get a copy of RFC1118, *Hitchhiker's Guide to the Internet*.

```
ftp> get RFC:RFC1118.TXT
200 Port 13.62 at host 130.180.4.1 accepted.
150 ASCII retrieve of TS:<RFC>RFC1118.TXT.1 (25 pages)
started.
226 Transfer completed. 62792 (8) bytes transferred.
local: RFC:RFC1118.TXT remote: RFC:RFC1118.TXT
62792 bytes received in 43 seconds (1.4 Kbytes/s)
ftp>
```

Logging Out of The Remote Host

To log out of the remote host, type quit, as follows:

```
ftp> quit
221 QUIT command received. Goodbye.
myhost>
```

Obtaining RFCs Through The NIC Automatic Mail Service

To obtain RFCs from the Network Information Center (NIC) through an automatic mail service, send requests to the following address:

```
service@nic.ddn.mil
```

Include the RFC file name in the subject field of the message. For example:

```
Subject: SEND RFC:RFCnnnn.TYPE
```

"nnnn" is the number of the RFC you want

"type" is either .PS for the postscript version, or .TXT for the ASCII text version

Obtaining RFCs by Telephone Request

To request an RFC by telephone, call the Network Information Center at 1-800-235-3155.

Glossary

This section defines terms used in this manual.

access control information

On DECnet networks, information specified to gain access to a VMS system, primarily a user name and password.

adaptive routing

A type of network routing in which the path to a destination node is determined according to changing conditions on the network.

addressing

On internet networks, the way the network defines the locations of hosts.

address mask

A bit mask used to select bits from an internet address for subnet addressing. The mask is 32 bits long and selects the network portion of the internet address and one or more bits of the local portion. Also called a network mask, netmask, or subnet mask.

Address Resolution Protocol (ARP)

A protocol on internet networks that allows hosts to determine the physical network addresses of other hosts on the same network.

area

A group of nodes on a DECnet network that can run independently as a subnetwork.

area router

Same as a level 2 router.

ARP

See *Address Resolution Protocol (ARP)*.

asynchronous connections

Connections that use asynchronous transmission, a mode of data transmission in which the time intervals between transmitted characters may be of unequal length. Asynchronous transmission most commonly occurs over terminal lines.

autonomous system

On internet networks, a collection of routers and networks that fall under one administrative entity and cooperate closely to propagate network reachability and routing information among themselves.

babbler

A device experiencing hardware problems that sends large amounts of data on the network, adversely affecting the general operation of the network. Also known as a screamer.

Berkeley Internet Name Domain (BIND)

A name service available on internet networks.

BIND

See *Berkeley Internet Name Domain*.

bouncing circuit

On DECnet networks, a circuit that continually alternates between circuit states.

branch

Same as subdomain.

bridge

A device used to expand a local area network by forwarding frames between data link layers associated with two different kinds of physical link. Also called a data link relay or level 2 relay.

bus topology

A network configuration in which all computers are connected to a single cable that runs the length of the network. Every computer has access to every other computer on the network directly through the transmission media.

cache

The storage area for the ARP bindings most recently requested by a host on an internet network.

carrier sense multiple access with collision detection (CSMA/CD)

A link management method used on Ethernet networks. CSMA/CD allows multiple stations to access a transmission medium (multiple access) by listening until no signals are detected (carrier sense), then transmitting and checking to see if more than one signal is present (collision detection).

channel

A means of data transmission.

channel service unit (CSU)

A device used on digital circuits that performs signal conversion. See also *data service unit*.

circuit

Virtual communication path between nodes or DTEs. Circuits operate over physical lines and are the media on which all I/O occurs.

circuit cost

On DECnet networks, a number that the network manager assigns to a circuit between two nodes.

circuit loopback test

A circuit-level loopback test that uses a loopback connector or a modem. See also *controller loopback* and *software loopback tests*.

circuit quality

A measure of the effectiveness of a circuit; the percentage of data transmitted that reaches its destination intact. Circuit quality is based on the throughput through the circuit, and data errors and timeouts recorded.

client

On an internet network, a host system that obtains naming information from a name server.

CSMA/CD

See *carrier sense multiple access with collision detection*.

common carrier

An organization that offers standard and consistent communications services within a country. For example, Western Union or American Telephone and Telegraph (AT&T).

common internet address notation

On internet networks, the decimal form of the 32-bit internet address. Also called dotted decimal notation.

communications channel

The path through which signals are sent to computers.

configuration database

On DECnet networks, a database containing files that provide information about network components. Specifically, the files contain information about the local node, and all remote nodes, modules, circuits, lines, logging, and objects in the network. See also *permanent database* and *volatile database*.

controller loopback test

A circuit-level loopback test that uses a hardware loopback device. See also *circuit loopback* and *software loopback tests*.

count

The number of blocks to be sent during a loopback test.

data communications equipment (DCE)

The network equipment (such as modems and multiplexors) that establishes, maintains, and terminates a connection, and handles the signal conversion and coding between the data terminal equipment and the network.

data communications network

Two or more computers set up to exchange information or data. Also known as a computer network.

data service unit (DSU)

A device used on digital circuits that performs signal conversion. See also *channel service unit*.

data terminal equipment (DTE)

User's equipment (computer or terminal) connected to a DCE on a packet switching network for the purpose of sending and receiving data.

DCE

See *data communications equipment*.

DDCMP

See *Digital Data Communication Message Protocol*.

DELNI

A local network interconnect device used to group Ethernet-compatible devices (not terminals) together.

Digital Data Communication Message Protocol (DDCMP)

A data link layer protocol where message headers contain a character count indicating the size of the message. This permits the transmission of arbitrary character sequences in the message without concern that some may be interpreted as control characters.

DIGITAL Network Architecture (DNA)

A set of protocols that govern the interrelationship of the components that make up the DECnet software. The specific functional boundaries between DECnet software components are structured as a hierarchical set of layers. DNA specifies the functional layers in which DECnet software is arranged and the communications protocols through which the corresponding layers at different nodes communicate with each other.

direct routing

On internet networks, the process of directing an IP datagram from one host to another host on the same physical network.

disruptive loopback test

Any loopback test that disconnects the existing circuit before looping the data, either physically with a loopback connector or button, or through software. Disruptive loopback tests disable any services the circuit may have been providing.

DNA

See *DIGITAL Network Architecture*.

DTE

See *data terminal equipment*.

domain

The primary category of hosts in the internet domain name system. The Internet domain name system consists of seven primary categories of hosts including COM, EDU, GOV, MIL, NET, ORG, and ARPA.

domain name system

A tree-structured system for organizing host names for an entire internet.

dotted decimal notation

Same as common internet address notation.

end node

A DECnet node that can receive packets addressed to it and send packets to other nodes, but cannot route packets through from other nodes. Also called a nonrouting node.

External Gateway Protocol (EGP)

A type of routing protocol that allows individual networks to communicate with the Internet backbone.

Ethernet

A CSMA/CD system using coaxial cable, developed at Xerox Corporation's Palo Alto Research Center used for local communications networks.

executor node

The DECnet node at which an NCP command executes.

extended local area network

A collection of local area networks connected by bridges, and logically considered as one local area network.

FDDI

See *Fiber Distributed Data Interface*.

Fiber Distributed Data Interface (FDDI)

A 100-Mbps local area network (LAN) standard that uses a timed-token access method and allows up to 500 stations to be connected with a total fiber length of 200 kms.

fiber optics

A data transmission medium consisting of fine glass fibers.

hardware loopback test

A type of loopback test that loops data through specific hardware points using a hardware loopback connector inserted on the end of a cable.

hop

The distance between two directly connected nodes on a DECnet network. On an internet network, a measure of distance between two points on the network.

hop count

On an internet network, the number of networks a datagram travels on its way from the source host to the destination host. If the destination host for an IP datagram is on the same network as the source host, the hop count is one, because the datagram travels on only one network. If the destination host is on another network, the hop count is two, because the datagram travels from one network, through a router, to a second network.

host

A computer system on an internet network.

host number

The second of two parts of an internet address. The host number identifies a host on the internet.

ICMP

See *Internet Control Message Protocol*.

indirect routing

The process of directing an IP datagram from one host on an internet network to another host on a different physical network.

Interior Gateway Protocol (IGP)

A type of routing protocol used to route information for networks within an autonomous system on an internet network.

internet

A collection of connected networks using the internet protocol (IP).

Internet

A collection of computing networks consisting of participants from major research institutions, universities, and government labs, including the National Science Foundation (NSF) and the NFSnet regional organizations. The Internet is not a commercial product but rather, a large project in support of research. The Internet is also known as the TCP/IP Internet.

internet address

A unique 32-bit number that identifies a host's connection to an internet network. An internet address consists of a network number and a host number.

Internet Control Message Protocol (ICMP)

The protocol responsible for reporting network errors on internet networks. ICMP also provides route change information from IP routers to hosts. The route change information contains alternate routes to use when hosts on the route are unavailable.

Internet Protocol (IP)

On internet networks, the layer upon which the Transmission Control Protocol (TCP) and application services build. IP determines how data is sent from one host to another on an internet, and specifies the format of internet packets known as internet protocol (IP) datagrams.

internet protocol (IP) datagram

The basic unit of information passed across an internet. An IP datagram contains source and destination addresses as well as data.

internet protocol (IP) gateway

Same as internet protocol (IP) router.

internet protocol (IP) router

A host that connects to two or more internet networks. The IP router knows how to reach all the hosts on the networks to which it is attached. Also called an internet protocol (IP) gateway.

internetwork

Same as internet.

IP

See *Internet Protocol (IP)*.

LAN

See *local area network*.

leaf

A host in the internet domain name system.

least cost path

The route to a DECnet node with the lowest sum of circuit costs.

length

Length (in bytes) of the blocks to be sent during a loopback test.

level 1 router

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

level 2 router

A DECnet node that 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 component that provides a distinct physical data path.

load host node

Any node that provides downline loading and upline dumping for other systems.

local node

The node at which you are physically located.

local area network (LAN)

A data communications system that operates over a limited physical distance, offering high-speed communications channels optimized for connecting information-processing equipment. Also known as a LAN.

Local Area Transport (LAT)

A communications protocol that the VMS operating system uses within a local area network to communicate with terminal servers.

logical link

A carrier of a single stream of full-duplex traffic between two user-level processes.

logical location

The functional interconnections between devices on a network. See also *physical location*.

management agent

A component of the Simple Network Management Protocol (SNMP) that resides on the network elements. The management agent performs network management functions requested by the network management stations that monitor and control network elements.

MAU

See *media access unit*.

media access unit (MAU)

A data communication connection used on LANs to provide an interface to the network for a computer.

mesh topology

A network configuration that uses multiple paths to each computer system to ensure that no single point of failure exists on the network.

modem

An acronym for modulator-demodulator. A device that changes digital signals into analog signals for transmission over long distances, and changes received analog signals to digital signals for use by electronic data equipment. Also, a device that converts computer signals to signals that can be sent over a telephone line.

name server

Software that maintains naming information for network entities and performs name-to-address and address-to-name translations. Sometimes the name server software runs on a dedicated processor, and the processor itself is also known as the name server.

naming

The method for identifying hosts on an internet network.

netmask

Same as address mask.

network management stations

A component of the Simple Network Management Protocol (SNMP) that monitors and controls network elements.

network elements

Devices such as hosts, gateways, and terminal servers. These devices contain Simple Network Management Protocol (SNMP) management agents, which perform network management functions requested by the network management stations.

Network Information Center (NIC)

The central authority for the Internet.

network mask

Same as address mask.

network number

The part of an internet address that encodes the network address and the network class of a host.

network problem

Any failure of hardware, software, or communications media that prevents users from accessing network resources.

NIC

See *Network Information Center*.

node

A computer system on a DECnet network.

nondisruptive loopback test

Any loopback test that does not disconnect the circuit before looping the data.

nonrouting node

An end node on a DECnet network.

object

A DECnet-VAX process that receives a logical link request. It performs a specific network function (a nonzero object such as FAL or NML), or is a user-defined image for a special-purpose application (a zero-numbered object).

octet

An 8-bit byte. Also, an 8-bit field of the 32-bit internet address.

packet rate

The number of packets transmitted per second.

partitioned area

An area in a DECnet network that, due to configuration weaknesses or multiple failures on the network, is unreachable to some, but not all other nodes and areas on the network.

path

The route a packet takes from a source to a destination.

path cost

The sum of the circuit costs along a path between two DECnet nodes. See also *circuit cost*.

path length

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

peer-to-peer network

A network, such as a DECnet network, in which all systems participate as peers or equals.

permanent database

Files containing information about network management components. See also *volatile database* and *configuration database*.

physical location

The place a device is stationed on a network. See also *logical location*.

point-to-point connections

A circuit that connects two nodes, operating over a single line.

PID

See *process identification*.

protocol

A set of messages with specific formats and the rules for exchanging the messages that governs the operation of a communications link.

process identification (PID)

A 32-bit binary value that uniquely identifies a process. Each process has a process identification and a process name. Also called PID.

Requests for Comments (RFC)

A series of notes that contain surveys, measurements, ideas, techniques, and observations, as well as proposed and accepted Internet protocol standards.

remote node

Any node in the network other than the node at which you are physically located.

resolver

Client software that queries name servers for the naming information requested.

RFC

See *Requests for Comments (RFC)*.

ring topology

A network configuration that consists of computers arranged in a closed loop.

root

The top-most level of the internet domain name system. The root level encompasses seven primary categories of internet hosts known as domains, including COM, EDU, GOV, MIL, NET, ORG, and ARPA.

root server

The top-level name servers for the internet domain name system.

routed

Route Daemon. The program that implements the Routing Information Protocol (RIP) on internet networks.

router

An intervening node along the data path between two nodes that forwards the data received from the source node to the destination node.

routing

The process of directing a data message from a source node or host to a destination node or host.

Routing Information Protocol (RIP)

An Interior Gateway Protocol (IGP) that specifies how routing information passes between routers in an autonomous system. The program that implements RIP is routed.

routing overhead

The amount of traffic generated to maintain the adaptive routing tables.

routing tables

The tables that hosts maintain to keep track of the best routes to other networks in an internet. For each destination network, routing tables list the network number, and the internet address to use when sending datagrams to that network.

screamer

See *babbler*.

service node

A VMS host system that is part of a LAT configuration.

Simple Network Management Protocol

An interim standard for network management in TCP/IP networks.

software loopback test

A circuit-level loopback test that loops data using the software capabilities of the system. See also *circuit loopback* and *controller loopback tests*.

SNMP

See *Simple Network Management Protocol*.

software loopback test

A type of loopback test that loops data through various DNA layers using NCP commands.

star topology

A network configuration in which each computer in the network is connected to a central computer through a data circuit.

subdomain

In the internet domain name system, the groups that participate in a domain. Also called zones or branches.

subnet

Each physical network that shares a network address with other physical networks on an internet. Because a single network address identifies all of the physical networks in a subnet collection, networks outside of the collection see only one network, where in fact there are multiple physical networks.

subnet address

A modified form of an internet address for internets that use subnets. A subnet address divides the 32-bit address into a network part and a local part. The network part identifies the internet network address. The local part identifies the physical network (subnet) and host number.

subnet mask

Same as address mask.

synchronous connections

Connections that use synchronous transmission, a mode of data transmission in which the time of occurrence of each signal representing a bit is related to a fixed time frame.

system user authorization file (SYSUAF)

A VAX/VMS file containing an entry for every user that the system manager authorizes to gain access to the system. Each entry identifies the user name, password, default account, user identification code, quotas, limits, and privileges assigned to individuals who use the system.

SYSUAF

See *system user authorization file*.

TCP/IP

Communications software used on internet networks, based loosely on the internet protocol model. TCP/IP consists of two main protocols: Transmission Control Protocol (TCP) and Internet Protocol (IP).

TCP/IP Internet

Same as the Internet.

terminal server

A device used to connect multiple terminals to service nodes in a network.

token ring topology

A network configuration that consists of computers arranged in a closed loop.

topology

The physical and logical location of components in a network.

transit congestion loss

The state where circuits or routers or both receive too many packets to be processed at one time. When this occurs, the packets are discarded, and the transport layer of the system that originated the packets retransmits the packets.

Transmission Control Protocol (TCP)

A host-to-host protocol for reliable communication in an internet.

tree topology

A network configuration in which a central, controlling computer manages the various groups of computers on the network.

User Datagram Protocol (UDP)

The internet standard protocol that allows an application program on one host to send a datagram to an application program on another host. UDP provides unreliable, connectionless delivery service using IP to transport messages among hosts.

volatile database

A memory image that contains information about network management components. Information in the volatile database is lost when the system reboots. See also *permanent database*.

WAN

See *wide area network*.

wide area network (WAN)

A public or private communications system, such as DECnet, that provides transmission to a wide geographical area.

Yellow Pages (YP)

A name server used on internet networks.

zone

Same as a subdomain.

Index

A

Aborted service request, 5-4
Access control rejected, 5-23
Accounting log file, 4-3
ACCOUNTNG.DAT file, 4-3
Adaptive routing, 1-10
Address mask, 1-28
Address Resolution Protocol (ARP), 1-39
Adjacency rejected, 5-8
Adjacency rejected/adjacency up, 5-8
Adjacency up, 5-8
ALIAS MAXIMUM LINKS
 guidelines for setting, 5-49
Architectures, 2-2
 Ethernet 802.3, 1-8
 using to solve DECnet problems, 2-2
 using to solve TCP/IP problems, 2-3
ARP
 See Address Resolution Protocol
arp command, 3-2
 using for direct routing problems, 3-2
ARP storm
 causes, 5-16
Asynchronous connection
 dynamic, 1-14
 static, 1-14
Asynchronous DECnet problems, 5-10
AUTHORIZE
 See Authorize utility
Authorize utility, 3-3

B

Babbling device, 5-14
Baseband Ethernet
 configuration guidelines, A-1
Berkeley Internet Name Domain (BIND), 1-32
BIND
 See Berkeley Internet Name Domain
Bouncing circuit, 5-8
Breakout boxes, 3-5
 checking modem signals with, 5-41
Bridges, 1-13
 configuration guidelines, A-3
 LAN Bridge switch settings, 5-57

Bridges (cont'd)
 METROWAVE, 1-13
Broadcast storm, 5-16

C

Cable, 1-8
 Ethernet, 1-7
Circuit, 1-8
 bouncing, 5-8
 cost, 1-9
Circuit-level loopback tests, 4-13
Circuit quality formula, 4-21
Circuit state problems, 5-20
Cleaning up after troubleshooting, 2-20
Cluster alias
 as cause of intermittent errors, 2-14
Common internet address notation, 1-21
Communications channel, 1-2
Computing system, 1-9
Configuration errors
 defined, 2-15
Configuration guidelines
 baseband Ethernet, A-1
 bridges, A-3
 DELNI local area interconnect, A-2
 DEMPR ThinWire multiport repeaters, A-2
 DEREP Ethernet repeaters, A-2
 fiber-optic cable, A-2
 ThinWire Ethernet, A-1
Configuring network devices on TCP/IP networks,
 5-43, 5-88
Connect failed, access control rejected, 5-23
Connect failed, unrecognized object, 5-29
Connection
 point-to-point, 1-14
Connection timed out, 5-31
 example, 2-6
 using ping command for solving, 3-23
Controller loopback tests, 4-13, 4-14
Cross-category problems
 circuit state problems, 5-20
 connection timed out, 5-31
 dialup, 5-38
 host is unreachable, 5-42
 Network is unreachable, 5-87
 remote node is not currently reachable, 5-109

CSMA/CD, 1-8

D

DECelms

Data communications equipment (DCE), 1-1

Data communications media, 1-2

Data communications network

 components of, 1-1

 defined, 1-1

Data link layer, 1-11

Data set hangup error, 5-68

Data terminal equipment (DTE), 1-1

DCE

 See data communications equipment

DDCMP

 See DIGITAL Data Communications Message Protocol

DEC Extended LAN Management Software (DECelms), 3-5

DECmcc Management Station for ULTRIX, 3-7

DECnet

 architecture, 1-10

 area, 1-7

 layers, 1-11

 node, 1-10

 protocol, 1-10

 structure, 1-10

DECnet area

 maximum number of nodes, 5-98

 maximum number per DECnet network, 5-98

DECnet-DOS software

 in a DECnet network, 1-12

DECnet errors during network startup, 5-51

DECnet events, 4-23

 logging facility, 4-23

DECnet layers

 data link, 1-11

 end communications (or transport), 1-11

 network (or routing), 1-11

 network application, 1-11

 physical, 1-11

 routing (or network), 1-11

 session control, 1-11

 transport (or end communications), 1-11

 user, 1-11

DECnet Monitor, 3-18

DECnet networks

 areas, 1-7

 communication, 1-8

 configuration, 1-7, 1-12

 data transmission media, 1-7

 environments, 1-12

 local area network, 1-13

 maximum number of areas, 5-98

 multiple-area networks, 1-7

 peer-to-peer communication, 1-6

DECnet networks (cont'd)

 required skills for network troubleshooting, xiii

 subnetworks, 1-7

 topologies, 1-7

DECnet node

 relocation, 1-7

DECnet Phase IV addresses, 5-35

DECnet problems

 due to Ethernet address modification, 5-35

 using architectural information to solve, 2-2

DECnet router

 area, 1-9

 level 1, 1-9

 level 2, 1-9

DECnet routing, 1-9

DECnet-RSX software, 1-12

DECnet/SNA gateway, 1-12, 1-15

DECnet-ULTRIX software, 1-12

DECnet-VAX

 adaptive routing, 1-10

 node, 1-12

 object, 1-9

DELNI local network interconnect, 1-13

 configuration guidelines, A-2

DEREP Ethernet repeaters

 configuration guidelines, A-2

Device not mounted, 5-36

Dialup problems, 5-38

DIGITAL Data Communications Message Protocol (DDCMP), 1-14

 asynchronous connection, 1-14

 synchronous connection, 1-14

Direct routing, 1-33

Disruptive loopback tests, 4-8

Distributed processing, 1-7

DNA (DIGITAL Network Architecture), 1-10

 layers, 1-10, 1-11

 protocols, 1-10

 specifications, 1-10

Domain name system, 1-29

 domain, 1-29

 leaf, 1-29

 root, 1-29

 subdomain, 1-29

Dotted decimal notation, 1-21

DTE

 See data terminal equipment

E

Echo tests

 using the ping command, 3-22, 4-17

EGP

 See External Gateway Protocol

End communications (or transport) layer, 1-11

End node, 1-9

ERRLOG.SYS file, 4-2
Errors
 displaying LAT error messages, 5-66
 displaying VMS errors, 4-2
Errors while starting DECnet
 See invalid parameter value
/etc/hosts file, 1-32
/etc/ifconfig command, 5-43, 5-88
/etc/password file, 5-78
/etc/rc.local file, 5-43, 5-88
/etc/resolv.conf file, 5-121
/etc/svc.conf file, 5-120
/etc/svcorder file, 5-120
Ethernet
 802.3, 1-8
 cable, 1-7, 1-8
 cabling rules, 2-18, A-1
 channel, 1-7
 concepts, 2-18
 configuration guidelines, A-1
 data transmission rate, 1-8
 extended, 1-8
 line statistics formula, 4-22
 maximum number of nodes, 1-8
 maximum number of taps, 1-8
 multi-access device, 1-7
 protocols, 2-18
 ThinWire, 1-8
Ethernet addresses, 5-35
 Phase IV DECnet's modification of, 5-35
ETHERnim, 3-20
Event logging
 enabling, 4-1
Extended Ethernet, 1-8
External Gateway Protocol (EGP), 1-37

F

Fault isolation
 defined, 2-15
Fiber-optic cable
 configuration guidelines, A-2
Fiber-optic link, 1-7
Formulas
 circuit quality, 4-21
 Ethernet line statistics, 4-22
 for understanding counters, 4-20
 packet rate, 4-20
 retransmissions, 4-23
 routing overhead, 4-23
 transit congestion loss, 4-22
ftp
 See Internet File Transfer Protocol

G

Gateway, 1-15
 DECnet, 1-12
 DECnet/SNA, 1-12, 1-15
Gateway, IP
 See IP router

H

Hard errors
 defined, 2-13
Hardware errors
 defined, 2-14
 gathering information about, 4-2, 4-18
Hardware loopback device, 4-13
Hardware loopback tests, 4-8
Historical performance data
 maintaining, 2-4
Hop
 DECnet network, 1-9
 internet network, 1-37
Hop count
 internet, defined, 1-37
Host
 relocation on internet, 1-18
Host is unreachable, 5-42
 using ping command for solving, 3-23
Host problems
 defined, 2-12
Host unknown
 See Unknown host
Host unreachable
 example, 2-6

I

ICMP
 See Internet Control Message Protocol
IGP
 See Interior Gateway Protocol
Inaccurate ARP entries
 how to correct, 5-35
Inconsistent errors
 defined, 2-13
Indirect routing, 1-33
 process, 1-33
Information gathering
 on hardware errors, 4-2, 4-18
 using event logging, 4-1
 using log files, 4-1
 using NCP counters, 4-18
 using SNMP tools, 4-23
 using the ACCOUNTNG.DAT file, 4-3
 using the DECnet event logging facility, 4-23
 using the ERRLOG.SYS file, 4-2

- Information gathering (cont'd)
 - using the NETSERVER.LOG file, 4-2
 - using the OPERATOR.LOG file, 4-1
 - using the ULTRIX error log files, 4-3
- Insufficient resources at remote node, 5-47
- Interior Gateway Protocol (IGP), 1-37
- Intermittent errors
 - cluster configuration as cause, 2-14
 - defined, 2-14
 - due to threshold values being reached, 2-14
 - effect on performance, 2-4
- Internet
 - defined, 1-17
 - TCP/IP, 1-19
- Internet addresses, 1-20
 - common notation, 1-21
 - decimal ranges, 1-24
 - dotted decimal notation, 1-21
 - fields, 1-21
 - host number, 1-21
 - modifications for subnets, 1-27
 - network number, 1-21
 - octets, 1-21
 - parts, 1-21
 - translating to Ethernet addresses, 3-3
- Internet Control Message Protocol (ICMP), 1-37
- Internet File Transfer Protocol (ftp), B-1
- Internet name domains, 1-30
- Internet network classes
 - size of, 1-23
- Internet Protocol (IP), 1-19, 1-36
 - RFC 791, 1-20
- Internet Standard Subnetting Procedure
 - RFC 950, 1-28
- Internet-to-Ethernet address translation tables
 - displaying and modifying using arp command, 3-2
- Interoperability errors
 - defined, 2-15
- Invalid login information at remote node
 - using Authorize utility to solve, 3-4
- Invalid parameter value, 5-51
- IP
 - See Internet Protocol
- IP addressing
 - defined, 1-20
- IP datagrams, 1-20
- IP gateway
 - See IP router
- IP hosts
 - testing reachability with the ping command, 3-22, 4-17
- IP naming
 - defined, 1-20
- IP router, 1-18

- IP routing
 - defined, 1-20
 - direct, 1-33
 - indirect, 1-33
- IP routing protocols
 - External Gateway Protocol, 1-37
 - Interior Gateway Protocol, 1-37
 - Routing Information Protocol, 1-37
- IP routing tables, 1-32
- Isolating the source of a problem, 2-15
 - to the host level, 2-16
 - to the LAN level, 2-18
 - to the node level, 2-16
 - to the WAN level, 2-18

L

- LAN
 - See local area network
- LAN Bridge cannot downline load, 5-53
 - using LAN Traffic Monitor to solve, 3-9
- LAN Bridge switch settings, 5-57
- LAN problems
 - aborted service request, 5-4
 - adjacency rejected/adjacency up, 5-8
 - babbling device, 5-14
 - bouncing circuit, 5-8
 - broadcast storm, 5-16
 - defined, 2-12
 - LAN Bridge cannot downline load, 5-53
 - LAN segment communication problem, 5-58
 - LAT port hung, 5-62
 - LAT print queue problems, 5-65
 - terminal server connection failures, 5-114
- LAN segment communication problem, 5-58
- LAN Traffic Monitor (LTM), 3-8
- LAT (local area transport)
 - protocol, 1-13
- LAT Control Program (LATCP), 3-10
 - running, 3-10
- LATCP
 - See LAT Control Program, 3-10
- LAT error messages
 - displaying, 5-66
- LAT port hung, 5-62
- LAT print queue
 - fixing a paused queue, 5-68
 - fixing a stalled queue, 5-67
 - fixing job retained on error problems, 5-72
 - fixing when printing incorrectly, 5-72
 - starting a stopped queue, 5-67
- LAT print queue problems, 5-65
- LCP
 - See LAT Control Program
- Leaf
 - in the domain name system, 1-29

- Least cost path, 1-10
- Level 1 router (DECnet), 1-9
- Level 2 router (DECnet), 1-9
- Line, 1-8
 - dedicated, 1-7, 1-14
 - dialup, 1-7, 1-14
 - terminal, 1-14
- Line synchronization lost, 5-74
 - using NMCC/VAX ETHERnim to solve, 3-22
- Link
 - fiber-optic, 1-7
 - microwave, 1-7, 1-8
 - satellite, 1-7, 1-8
- Local area network (LAN), 1-5
 - bridge, 1-13
 - configuration, 1-13
 - data communications connections, 1-1
 - DECnet, 1-13
- Local loopback tests, 4-12
- Local network interconnect device
 - DELNI, 1-13
- Local-to-local loopback tests, 4-11
- Local-to-remote loopback tests, 4-11
- Location
 - logical, 1-2
 - physical, 1-2
- Log files, 4-1
 - ACCOUNTNG.DAT, 4-3
 - ERRLOG.SYS, 4-2
 - NETSERVER.LOG, 4-2
 - OPERATOR.LOG, 4-1
- Logical link, 1-9
- Logical location, 1-2
- Login incorrect, 5-78
- Login information invalid, 2-3, 5-80
 - architectural layers affected, 2-3
- Loopback connector, 4-15
- Loopback tests
 - disruptive, 4-8
 - failed NCP tests, 4-8
 - modem, 4-15
 - nondisruptive, 4-9
 - remote, 4-9
 - results of NCP tests, 4-8
 - successful NCP tests, 4-8
 - to a remote node, 4-9
 - using a loopback connector, 4-15
 - using NCP, 4-6
 - using the ping command, 3-22
- Loop node name, 4-10
- LTM
 - See LAN Traffic Monitor

M

- Maps, 2-1
 - creating manually, 2-2
 - creating using software tools, 2-1
 - generating manually, 4-3
 - tools for creating, 2-1
 - use during troubleshooting, 2-1
- MAU
 - See media access units
- MAXIMUM LINKS
 - guidelines for setting, 5-49
- Media access units (MAU), 1-1
- METROWAVE bridge, 1-13
- Microwave link, 1-7, 1-8
- Modem, 1-7, 1-14
- Modem loopback tests, 4-13, 4-15
 - example of, 4-15
 - flowchart of, 4-16
 - for problems between a modem and a circuit, 4-15
- Modem signals
 - checking with a breakout box, 5-41
- Multi-access device, 1-7
- Multiple-area networks, 1-7

N

- NAKS
 - See Negative acknowledgements
- Name server, 1-31
 - Berkeley Internet Name Domain (BIND), 1-32
 - client, 1-31
 - resolver, 1-31
 - root servers, 1-32
 - Yellow Pages (YP), 1-32
- Naming information
 - in the /etc/hosts file, 1-32
- NCP
 - See Network Control Program
- NCP loopback tests, 4-6
 - circuit-level, 4-13
 - controller, 4-13, 4-14
 - failed, 4-8
 - hardware, 4-8
 - local, 4-12
 - local-to-local, 4-11
 - local-to-remote, 4-11
 - modem, 4-15
 - node-level, 4-9
 - remote loopback, 4-9
 - results of, 4-8
 - software, 4-8, 4-13, 4-14
 - successful, 4-8
 - types, 4-8
 - using a loop node name, 4-10

- Negative acknowledgements, 4-8
- Netmask, 1-28
- NETPROXY.DAT file, 3-3
- NETSERVER.LOG file, 4-2
- Netserver log file, 4-2
- netstat command, 3-12
 - tracing a routing path with, 4-5
- Network
 - architectures, 2-2
 - error sources, 2-14
 - gateway, 1-15
 - interconnect products, 1-12
 - maps, 2-1, 4-3
 - normal operation, 2-1
 - packet-switching, 1-12, 1-15
 - peer-to-peer, 1-6
 - performance, 2-4
 - topology, 1-2
 - types of errors, 2-13
 - understanding typical use, 2-4
 - wide area network, 1-13
- Network (or routing) layer, 1-11
- Network application layer, 1-11
- Network configurations, 1-2
 - bus, 1-4
 - mesh, 1-4
 - ring or loop, 1-3
 - star, 1-3
 - tree, 1-5
- Network Control Program (NCP), 3-13
 - counters, 4-18
 - DEFINE and LIST commands, 3-15
 - modifying the permanent databases, 3-15
 - modifying the volatile databases, 3-15
 - performing operations on remote nodes, 3-15
 - SET and DEFINE commands, 5-3
 - SET and SHOW commands, 3-15
 - tracing a routing path with, 4-4
- Network data structures
 - displaying with netstat command, 3-12
- Network devices
 - configuring on TCP/IP networks, 5-43, 5-88
- Network errors
 - configuration, 2-15
 - defined, 2-13
 - hard, 2-13
 - hardware, 2-14
 - inconsistent, 2-13
 - intermittent, 2-14
 - interoperability, 2-15
 - software, 2-15
 - sources, 2-14
 - transient, 2-14
 - types, 2-13
 - user, 2-14
- Network Information Center (NIC), 1-21
- Network is unreachable, 5-87
 - using ping command for solving, 3-23
- Network management and troubleshooting tools
 - summary chart, 3-1
- Network map
 - generating, 4-3
- Network mask, 1-28
- Network object, 1-9
 - DECnet-VAX system program, 1-9
 - MAIL, 1-9
- Network object unknown, 5-91
 - using Network Control Program to solve, 3-16
- Network partner exited, 5-94
- Network problems
 - cluster alias as cause, 2-14
 - defined, 2-9
 - detecting, 2-9
 - extent of disturbance on the network, 5-1
 - listed, 5-1
 - solving, 2-19
 - sources of errors, 2-14
 - understanding the extent of, 2-12
 - understanding the types of errors, 2-13
- Network proxy authorization file, 3-3
- Network tools
 - DECmcc Management Station for ULTRIX, 2-4
 - LAN Traffic Monitor, 2-4
 - Network Control Program (NCP), 2-4
 - NMCC/VAX ETHERnim, 2-4
 - NMCC/DECnet Monitor, 2-4
- Network topology, 1-2
 - bus, 1-4
 - generating information on, 4-3
 - mesh, 1-4
 - ring or loop, 1-3
 - star, 1-3
 - tree, 1-5
- NIC
 - See Network Information Center
- NMCC/VAX ETHERnim, 3-20
- NMCC/DECnet Monitor, 3-18
- Node, 1-8
 - end node, 1-9
 - maximum number per DECnet area, 5-98
 - routing node, 1-9
- Node-level loopback tests
 - for logical link operation, 4-9
 - for operation over specific circuit, 4-9
- Node out of range packet loss, 5-98
- Node problems
 - defined, 2-12
- Nondisruptive loopback tests, 4-9

O

- Object, network, 1-9
 - See network object
 - DECnet-VAX system program, 1-9
 - MAIL, 1-9
 - user-written program, 1-9
- OCC
 - See Office communications cabinet
- Office communications cabinet, 5-63
- OPCOM
 - See Operator communication manager
- OPERATOR.LOG file, 4-1
- Operator communication manager
 - running, 4-1
- Operator log file, 4-1

P

- Packet rate formula, 4-20
- Packet-switching network, 1-12, 1-15
- Partial routing update loss, 5-98
 - See node out of range packet loss
- Partitioned area, 5-102
 - using NMCC/DECnet Monitor to solve, 3-20
- Passwords
 - authority required for modifying, 5-3
 - modifying, 5-3
- Path
 - isolating problems on, 4-3
 - least cost, 1-10
 - routing, 1-9
- Paused LAT print queue
 - fixing, 5-68
- Peer-to-peer network, 1-6
- Performance, 2-4
 - maintaining historical data on, 2-4
 - threshold, 2-4
 - trend analysis, 2-4
 - usage peak, 2-4
- Performance data
 - maintaining, 2-4
- Performance problems
 - due to intermittent errors, 2-4
 - due to transient errors, 2-4
- Performance tools
 - DECmcc Management Station for ULTRIX, 2-4
 - LAN Traffic Monitor, 2-4
 - Network Control Program (NCP), 2-4
 - NMCC/VAX ETHERnim, 2-4
 - NMCC/DECnet Monitor, 2-4
- Permission denied, 5-106
- Personal computer
 - connection to a DECnet network, 1-12
 - in a DECnet network, 1-12

- Physical layer, 1-11
- Physical location, 1-2
- ping command, 3-22
 - for testing reachability on TCP/IP networks, 4-17
- Print queue
 - fixing a paused LAT queue, 5-68
 - fixing a stalled LAT queue, 5-67
 - fixing job retained on error problems, 5-72
 - fixing when printing incorrectly, 5-72
 - starting a stopped queue, 5-67
- Print queue jobs retained on error
 - fixing, 5-72
- Print queue states and conditions, 5-67
- Privileged accounts
 - required for troubleshooting, 5-3
- Problem detection, 2-9
- Problems between a modem and a circuit
 - isolating, 4-15
- Problem statement, 2-9
- Processes
 - communication with, 1-9
- PRO/DECnet software, 1-12
- Professional 300-series system
 - in a DECnet network, 1-12
- Protocol
 - communications, 1-10
 - DDCMP, 1-14
 - DNA, 1-10
 - LAT, 1-13
- Protocol analyzers, 3-23

R

- Reachability of IP hosts
 - testing with the ping command, 3-22, 4-17
- Remote loopback tests, 4-9
- Remote node is not currently reachable, 2-2, 5-109
 - architectural layers affected, 2-2
 - example, 2-6
 - using NMCC/DECnet Monitor to solve, 3-20
- Request for Comments (RFC), 1-18
 - A Simple Network Management Protocol, Number 1098, 1-40
 - copying from public directories, B-3
 - Internet Protocol Model, Number 791, 1-20
 - Internet Standard Subnetting Procedure, Number 950, 1-28
 - obtaining an index of, B-2
 - obtaining through automatic mail service, B-3
 - obtaining through ftp, B-1
 - obtaining through telephone request, B-3
 - RFC 1098, 1-40
 - RFC 791, 1-20
 - RFC 950, 1-28

- Requirements for troubleshooting, 5-3
 - DECnet skills, xiii
 - TCP/IP skills, xiii
 - ULTRIX privileges, xiv
 - VMS privileges, xiv
- Resource sharing, 1-7
- Retained on error
 - print queue message, 5-72
- Retransmissions formula, 4-23
- RFC
 - See Request for Comments
- Rights database file, 3-3
- RIGHTSLIST.DAT file, 3-3
- RIP
 - See Routing Information Protocol
- Router, 1-9
- Routing
 - adaptive, 1-10
 - DECnet, 1-9
 - IP direct, 1-33
 - IP indirect, 1-33
 - path, 1-9
 - path cost, 1-9
 - path length, 1-9
- Routing (or network) layer, 1-11
- Routing Information Protocol (RIP), 1-37
- Routing overhead formula, 4-23
- Routing path
 - tracing, 4-3
 - tracing with NCP, 4-4
 - tracing with netstat, 4-5
 - tracing with the traceroute command, 3-27
- Routing tables, 1-32
- RSX system
 - in a DECnet network, 1-12

S

- Satellite equipment room (SER), 5-63
- Satellite link, 1-7, 1-8
- Screaming node
 - See Babbling device
- SER
 - See satellite equipment room (SER)
- Server
 - See terminal server
- Session control layer, 1-11
- Simple Network Management Protocol (SNMP), 1-40
 - RFC 1098, 1-40
- SNMP
 - See Simple Network Management Protocol
- SNMP tools
 - gathering information about ULTRIX counters with, 4-23

- Software errors
 - defined, 2-15
- Software loopback tests, 4-8, 4-13, 4-14
- Stalled LAT print queue
 - fixing, 5-67
- Stopped LAT print queue
 - starting, 5-67
- Structured problem-solving, 2-1
- Subdomain
 - in the domain name system, 1-29
- Subnet, 1-24
- Subnet addresses, 1-27
 - local part, 1-27
 - modification of internet address, 1-24
 - network part, 1-27
- Subnet mask, 1-28
- Subnets, 1-27
 - address masks, 1-28
 - netmasks, 1-28
 - network masks, 1-28
 - subnet masks, 1-28
- syslog daemon, 3-24
- System messages
 - displaying the ULTRIX error log file, 3-30
 - recording for ULTRIX, 3-24
- System user authorization file, 3-3
- SYSUAF.DAT file, 3-3

T

- TCP
 - See Transmission Control Protocol
- TCP/IP, 1-19
- TCP/IP Internet, 1-18
- TCP/IP loopback tests
 - using localhost, 4-18
 - using the ping command, 4-17
 - using the software loopback interface, 4-18
- TCP/IP networks
 - required skills for troubleshooting, xiii
- TCP/IP problems
 - solving using architectural information, 2-3
- Telephone line, 1-8, 1-14
 - dialup, 1-7
 - leased, 1-7
- Terminal server, 1-13
 - commands, 5-66
 - managing using TSM or NCP, 5-66
- Terminal server connection failure
 - using TSM to solve, 3-27
- Terminal server connection failures, 5-114
- Terminal Server Manager Software (TSM), 3-26
- Terminal server problems
 - data set hangup error, 5-68
- Tests
 - circuit-level loopback, 4-13
 - controller loopback, 4-14
 - failed NCP loopback, 4-8

Tests (cont'd)

- local loopback, 4-12
- local-to-local loopback, 4-11
- local-to-remote loopback, 4-11
- loopback results using NCP, 4-8
- loopback using NCP, 4-6
- node-level loopback, 4-9
- remote loopback, 4-9
- software loopback, 4-14
- successful NCP loopback, 4-8
- TCP/IP host reachability, 4-17

ThinWire Ethernet, 1-8

- configuration guidelines, A-1

ThinWire multiport repeaters

- configuration guidelines, A-2

Threshold, 2-4

Tools

- network management and troubleshooting
 - summary chart, 3-1

Topology, 1-2

- local area network, 1-5
- logical location, 1-2
- physical location, 1-2
- wide area network, 1-6

traceroute command, 3-27

Tracing a routing path, 4-3

- using NCP, 4-4
- using netstat, 4-5
- using the traceroute command, 3-27

Transient errors

- defined, 2-14
- due to threshold values being reached, 2-14
- effect on performance, 2-4
- occurring at peak usage times, 2-14

Transit congestion loss formula, 4-22

Transmission Control Protocol (TCP), 1-19, 1-38

Transport (or end communications) layer, 1-11

Transport protocols

- Transmission Control Protocol (TCP), 1-38
- User Datagram Protocol (UDP), 1-39

Trend analysis, 2-4

Troubleshooting

- authority required for modifying passwords, 5-3
- using privileged accounts, 5-3
- using the SET and DEFINE commands in NCP, 5-3

Troubleshooting methodology, 2-1

- analyzing, interpreting, and classifying information, 2-11
- cleaning up after troubleshooting, 2-20
- detecting a problem, 2-9
- documenting the problem and solution, 2-20
- gathering information, 2-9
- isolating the source of the problem, 2-15
- obtaining a problem statement, 2-9
- overview, 2-5
- solving the problem, 2-19

Troubleshooting methodology (cont'd)

- steps, 2-5
- verifying the solution, 2-20

TSM

- See Terminal Server Manager Software

U

UDP

- See User Datagram Protocol

uerf command, 3-30

ULTRIX commands

- arp, 3-2
- netstat, 3-12, 4-5
- ping, 3-22
- traceroute, 3-27
- uerf, 3-30

ULTRIX counters, 4-23

- gathering information using SNMP tools, 4-23

ULTRIX files

- /etc/password, 5-78
- /etc/rc.local, 5-43, 5-88
- /etc/resolv.conf, 5-121
- /etc/svc.conf, 5-120
- /etc/svcorder, 5-120

ULTRIX host problems

- connect failed, access control rejected, 5-23
- connect failed, unrecognized object, 5-29
- connection timed out, 5-31
- login incorrect, 5-78
- permission denied, 5-106
- unknown host, 5-119

ULTRIX system

- in a DECnet network, 1-12

ULTRIX system messages

- displaying the error log file, 3-30
- recording, 3-24

ULTRIX systems

- required privileges for troubleshooting, xiv

ULTRIX troubleshooting tools

- arp command, 3-2
- DECmcc Management Station for ULTRIX, 3-7
- lcp, 3-10
- netstat command, 3-12
- Network Control Program (ncp), 3-13
- ping command, 3-22
- syslog daemon, 3-24
- traceroute command, 3-27
- uerf command, 3-30

Unknown host, 5-119

Unrecognized object, 5-29

Usage peak, 2-4

User Datagram Protocol (UDP), 1-39

User errors

- defined, 2-14

User layer, 1-11

V

VAX PSI software, 1-12, 1-15

Verification reject, 5-123

VMS error log file, 4-2

- displaying information from, 4-2
- generating a complete report, 4-2

VMS log files

ACCOUNTING.DAT, 4-3

NETSERVER.LOG, 4-2

OPCOM.LOG, 4-1

VMS node problems

- device not mounted, 5-36
- insufficient resources at remote node, 5-47
- invalid parameter value, 5-51
- line synchronization lost, 5-74
- login information invalid, 5-80
- network object unknown, 5-91
- network partner exited, 5-94
- node out of range packet loss, 5-98
- partial routing update loss, 5-98
- Verification reject, 5-123

VMS system

- communication with non-Digital systems, 1-12
- communication with other VMS systems, 1-12

VMS systems

- required privileges for troubleshooting, xiv

VMS troubleshooting tools

- Authorize utility, 3-3

DEC Extended LAN Management Software,
3-5

LAN Traffic Monitor (LTM), 3-8

LATCP, 3-10

Network Control Program (NCP), 3-13

NMCC/VAX ETHERnim, 3-20

NMCC/DECnet Monitor, 3-18

Terminal Server Manager Software, 3-26

W

WAN

- See Wide area network, 1-6

WAN problems

- Asynchronous DECnet problems, 5-10
- defined, 2-13
- partitioned area, 5-102

Wide area network (WAN), 1-6, 1-13

- configuration, 1-13
- data communications connections, 1-1

X

X.25 communications server, 1-12

X.25 packet-switching network, 1-15

Y

Yellow Pages (YP), 1-32

YP

- See Yellow Pages