

Compaq TCP/IP Services for OpenVMS

Tuning and Troubleshooting

Order Number: AA-RN1VA-TE

January 2001

This manual describes how to troubleshoot Compaq TCP/IP Services for OpenVMS and how to tune the performance of the product.

Revision Information:	This is a new manual.
Operating System:	OpenVMS Alpha Versions 7.1 and 7.2-1 OpenVMS VAX Versions 7.1 and 7.2
Software Version:	Compaq TCP/IP Services for OpenVMS Version 5.1

**Compaq Computer Corporation
Houston, Texas**

© 2001 Compaq Computer Corporation

COMPAQ, VAX, VMS, and the Compaq logo Registered in U.S. Patent and Trademark Office.

OpenVMS and Tru64 are trademarks of Compaq Information Technologies Group, L.P. in the United States and other countries.

UNIX is a trademark of The Open Group.

All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

ZK6631

This document is available on CD-ROM.

This document was prepared using DECdocument, Version 3.3-1b.

Contents

Preface	v
1 Troubleshooting Techniques and Tools	
1.1 Using Symptoms to Identify a Problem	1-1
1.2 Isolating Problems	1-1
1.2.1 Testing Connectivity Between Network Hosts	1-3
1.2.1.1 Using ping on a Multihomed Host	1-5
1.2.2 Checking the Network Interface Parameters	1-5
1.2.3 Displaying and Modifying the Internet-to-Ethernet Translation Tables	1-6
1.2.4 Examining Network Statistics	1-6
1.2.5 Monitoring Network Traffic	1-7
1.2.6 Checking Name Server Operation	1-8
1.2.7 Checking the Route to a Remote Host	1-10
1.2.8 Checking the Routes Known to a Gateway	1-11
1.2.9 Determine Whether Network Services Are Available	1-11
1.2.9.1 Displaying the Service Database	1-12
1.2.9.2 Displaying Service Attributes	1-13
1.2.9.3 Verifying Process Privileges	1-14
1.2.9.4 Verifying Account Privileges	1-14
1.2.9.5 Looking for OPCOM Messages	1-15
2 Tuning Techniques	
2.1 Kernel Subsystems	2-1
2.1.1 Displaying Subsystems and Attributes	2-1
2.1.2 Testing Attribute Support	2-2
2.1.3 Displaying Attribute Values	2-2
2.1.4 Modifying Attribute Values	2-3
2.1.5 Modifying Attributes Temporarily	2-3
2.2 Modifying Kernel Subsystems	2-3
2.2.1 Modifying Socket Subsystem Attributes	2-4
2.2.1.1 Increasing the Maximum Number of Pending TCP Connections	2-5
2.2.1.2 Increasing the Minimum Number of Pending TCP Connections	2-5
2.2.1.3 Increasing the Maximum Size of a Socket Buffer	2-5
2.2.2 Modifying Internet Subsystem Attributes	2-6
2.2.2.1 Increasing the Size of a TCP Hash Table	2-6
2.2.2.2 Increasing the Number of TCP Hash Tables	2-6
2.2.2.3 Increasing the Size of the Kernel Interface Alias Table	2-7
2.2.2.4 Increasing the TCP Partial Connection Timeout Rate	2-7
2.2.2.5 Slowing TCP Retransmission Rate	2-8
2.2.2.6 Enabling the TCP Keepalive Function	2-9

2.2.2.7	Increasing the Timeout Rate for TCP Connection Context	2-9
2.2.2.8	Modifying the Range of Outgoing Connection Ports	2-10
2.2.2.9	Disabling Use of the PMTU Protocol	2-10
2.3	Monitoring Servers	2-11
2.3.1	Displaying Network Statistics	2-11
2.3.2	Displaying Socket Statistics	2-13
2.4	Tuning Server Applications	2-13
2.4.1	Configuring Memory for High Performance	2-13
2.4.2	Logging IP Addresses	2-14
2.5	Increasing the Auxiliary Server Connection Limit	2-14
2.6	Solving Performance Problems	2-14
2.6.1	Tuning Recommendations for a Primary Server	2-15

A Troubleshooting Tools Reference

arp	A-2
dig	A-5
ifconfig	A-11
ndc	A-16
netstat	A-18
nslookup	A-22
ping	A-27
route	A-31
sysconfig	A-35
TCPTRACE	A-38
traceroute	A-40

B Commonly Used UNIX Commands

Index

Tables

1	TCP/IP Services Documentation	vi
1-1	Diagnostic Tools	1-2
2-1	Network Tuning Guidelines	2-3
2-2	socket Subsystem Attributes	2-4
2-3	inet Subsystem Attributes	2-6
2-4	TCP Keepalive Options	2-9
A-1	dig Options	A-6
A-2	Options to the nslookup set Command	A-24
B-1	Commonly Used Commands	B-2

Preface

Compaq TCP/IP Services for OpenVMS is the Compaq implementation of the TCP/IP networking protocol suite and internet services for OpenVMS Alpha and OpenVMS VAX systems.

A layered software product, TCP/IP Services provides a comprehensive suite of functions and applications that support industry-standard protocols for heterogeneous network communications and resource sharing.

See the *Compaq TCP/IP Services for OpenVMS Installation and Configuration* manual for information about installing, configuring, and starting this product.

This manual provides system and network managers with information they need to identify and resolve problems. This manual is best used in conjunction with the *Compaq TCP/IP Services for OpenVMS Management* manual.

Intended Audience

This manual is for OpenVMS or UNIX system managers who are experienced in troubleshooting complex software products. This manual assumes a working knowledge of TCP/IP networking, TCP/IP terminology, and familiarity with the TCP/IP Services for OpenVMS product. If you are not familiar with the TCP/IP Services for OpenVMS product, please review all product documentation before attempting to resolve any problems.

Document Structure

This manual contains the following two chapters and appendix:

- Chapter 1 describes how to determine the cause of networking problems. It introduces some tools useful in monitoring and diagnosing these problems.
- Chapter 2 describes the TCP/IP subsystem attributes that you can adjust to improve network performance.
- Appendix A describes the tools available for isolating and resolving network problems.

Related Documents

Table 1 lists the documents available with this version of TCP/IP Services.

Table 1 TCP/IP Services Documentation

Manual	Contents
<i>DIGITAL TCP/IP Services for OpenVMS Concepts and Planning</i>	<p>This manual provides conceptual information about networking and the TCP/IP protocol including a description of the Compaq implementation of the Berkeley Internet Name Domain (BIND) service and the Network File System (NFS). It outlines general planning issues to consider before configuring your system to use the TCP/IP Services software.</p> <p>This manual also describes the manuals in the documentation set, provides a glossary of terms and acronyms for the TCP/IP Services software product, and documents how to contact the InterNIC Registration Service to register domains and access Requests for Comments (RFCs).</p>
<i>Compaq TCP/IP Services for OpenVMS Release Notes</i>	<p>This text file describes new features and changes to the software including installation, upgrade, configuration, and compatibility information. These notes also describe new and existing software problems and restrictions, and software and documentation corrections.</p> <p>Print this text file at the beginning of the installation procedure and read it before you install TCP/IP Services.</p>
<i>Compaq TCP/IP Services for OpenVMS Installation and Configuration</i>	<p>This manual explains how to install and configure the TCP/IP Services product.</p>
<i>DIGITAL TCP/IP Services for OpenVMS User's Guide</i>	<p>This manual describes how to use the applications available with TCP/IP Services such as remote file operations, email, TELNET, TN3270, and network printing. This manual explains how to use these services to communicate with systems on private internets or on the worldwide Internet.</p>
<i>Compaq TCP/IP Services for OpenVMS Management</i>	<p>This manual describes how to configure and manage the TCP/IP Services product.</p> <p>Use this manual with the <i>Compaq TCP/IP Services for OpenVMS Management Command Reference</i> manual.</p>
<i>Compaq TCP/IP Services for OpenVMS Management Command Reference</i>	<p>This manual describes the TCP/IP Services management commands.</p> <p>Use this manual with the <i>Compaq TCP/IP Services for OpenVMS Management</i> manual.</p>
<i>Compaq TCP/IP Services for OpenVMS Management Command Quick Reference Card</i>	<p>This reference card lists the TCP/IP management commands by component and describes the purpose of each command.</p>
<i>Compaq TCP/IP Services for OpenVMS UNIX Command Reference Card</i>	<p>This reference card contains information about commonly performed network management tasks and their corresponding TCP/IP management and Compaq <i>Tru64</i> UNIX command formats.</p>
<i>DIGITAL TCP/IP Services for OpenVMS ONC RPC Programming</i>	<p>This manual presents an overview of high-level programming using open network computing remote procedure calls (ONC RPCs). This manual also describes the RPC programming interface and how to use the RPCGEN protocol compiler to create applications.</p>

(continued on next page)

Table 1 (Cont.) TCP/IP Services Documentation

Manual	Contents
<i>Compaq TCP/IP Services for OpenVMS Sockets API and System Services Programming</i>	This manual describes how to use the Sockets API and OpenVMS system services to develop network applications.
<i>Compaq TCP/IP Services for OpenVMS SNMP Programming and Reference</i>	This manual describes the Simple Network Management Protocol (SNMP) and the SNMP application programming interface (eSNMP). It describes the subagents provided with TCP/IP Services, utilities provided for managing subagents, and how to build your own subagents.
<i>Compaq TCP/IP Services for OpenVMS Tuning and Troubleshooting</i>	This manual provides information about how to isolate the causes of network problems and how to tune the TCP/IP Services software for the best performance.
<i>Compaq TCP/IP Services for OpenVMS Guide to IPv6</i>	This manual describes the IPv6 environment, the roles of systems in this environment, the types and function of the different IPv6 addresses, and how to configure TCP/IP Services to access the 6bone network.

For additional information about Compaq *OpenVMS* products and services, access the Compaq website at the following location:

<http://www.openvms.compaq.com/>

If you are looking for a comprehensive overview of the TCP/IP protocol suite, you might find the following books useful:

- Comer, Douglas E. *Internetworking with TCP/IP Volume 1: Principles, Protocols, and Architecture*. Third edition. Englewood Cliffs, NJ: Prentice-Hall Inc., 1995.
- Stevens, W. Richard. *UNIX Network Programming Volume 1: Networking APIs: Sockets and XTI*. Second edition. Englewood Cliffs, NJ: Prentice-Hall Inc., 1999.
- Stevens, W. Richard. *UNIX Network Programming Volume 2: Interprocess Communications*. Second edition. Englewood Cliffs, NJ: Prentice-Hall Inc., 1999.

Reader's Comments

Compaq welcomes your comments on this manual. Please send comments to either of the following addresses:

Internet	openvmsdoc@compaq.com
Mail	Compaq Computer Corporation OSSG Documentation Group, ZKO3-4/U08 110 Spit Brook Rd. Nashua, NH 03062-2698

How to Order Additional Documentation

Visit the following World Wide Web address for information about how to order additional documentation:

<http://www.openvms.compaq.com/>

If you need help deciding which documentation best meets your needs, call 800-282-6672.

Conventions

The name TCP/IP Services means both:

- Compaq TCP/IP Services for OpenVMS Alpha
- Compaq TCP/IP Services for OpenVMS VAX

The following conventions are used in this manual. In addition, please note that all IP addresses are fictitious.

Ctrl/ <i>x</i>	A sequence such as Ctrl/ <i>x</i> indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.
PF1 <i>x</i>	A sequence such as PF1 <i>x</i> indicates that you must first press and release the key labeled PF1 and then press and release another key or a pointing device button.
Return	<p>In examples, a key name enclosed in a box indicates that you press a key on the keyboard. (In text, a key name is not enclosed in a box.)</p> <p>In the HTML version of this document, this convention appears as brackets, rather than a box.</p>
...	<p>A horizontal ellipsis in examples indicates one of the following possibilities:</p> <ul style="list-style-type: none">• Additional optional arguments in a statement have been omitted.• The preceding item or items can be repeated one or more times.• Additional parameters, values, or other information can be entered.
. . .	A vertical ellipsis indicates the omission of items from a code example or command format; the items are omitted because they are not important to the topic being discussed.
()	In command format descriptions, parentheses indicate that you must enclose choices in parentheses if you specify more than one.
[]	In command format descriptions, brackets indicate optional choices. You can choose one or more items or no items. Do not type the brackets on the command line. However, you must include the brackets in the syntax for OpenVMS directory specifications and for a substring specification in an assignment statement.
	In command format descriptions, vertical bars separate choices within brackets or braces. Within brackets, the choices are optional; within braces, at least one choice is required. Do not type the vertical bars on the command line.
{ }	In command format descriptions, braces indicate required choices; you must choose at least one of the items listed. Do not type the braces on the command line.
bold text	This typeface represents the introduction of a new term. It also represents the name of an argument, an attribute, or a reason.

<i>italic text</i>	Italic text indicates important information, complete titles of manuals, or variables. Variables include information that varies in system output (Internal error <i>number</i>), in command lines (/PRODUCER= <i>name</i>), and in command parameters in text (where <i>dd</i> represents the predefined code for the device type).
UPPERCASE TEXT	Uppercase text indicates a command, the name of a routine, the name of a file, or the abbreviation for a system privilege.
Monospace text	Monospace type indicates code examples and interactive screen displays. This typeface indicates UNIX system output or user input, commands, options, files, directories, utilities, hosts, and users. In the C programming language, this typeface identifies the following elements: keywords, the names of independently compiled external functions and files, syntax summaries, and references to variables or identifiers introduced in an example.
-	A hyphen at the end of a command format description, command line, or code line indicates that the command or statement continues on the following line.
numbers	All numbers in text are assumed to be decimal unless otherwise noted. Nondecimal radices—binary, octal, or hexadecimal—are explicitly indicated.

Troubleshooting Techniques and Tools

This chapter provides information that helps you identify symptoms, isolate problems, and take steps to resolve your network problem. This chapter also introduces the tools available to help you monitor and diagnose problems with your network software, devices, and interfaces.

1.1 Using Symptoms to Identify a Problem

The inability to reach remote hosts and networks is usually caused by one of the following:

- Physical connection failure
- Underlying transport failure (UDP, TCP, IP)
- Incorrectly configured routing, applications, or services such as BIND
- User error

1.2 Isolating Problems

The first step in problem isolation is to make sure that the TCP/IP Services product is started. This may seem like an obvious step, but it is frequently overlooked because error messages may not indicate the product is disabled. (Instead, the messages returned may be “invalid host” or something similar.) You may not have stopped the product, but someone else may have. To check whether the product is running, enter the following command:

```
$ SHOW DEVICE BG
Device          Device          Error
Name           Status         Count
BG0:           Mounted        0
BG1:           Mounted        0
BG5:           Mounted        0
BG6:           Mounted        0
BG7:           Mounted        0
BG8:           Mounted        0
```

If the command output shows only the BG0: device, then the product is stopped.

The second step is to reduce the problem to its basic components and to systematically identify what is and what is not working. Ask the following questions:

- Does the problem occur all the time or intermittently?
- Does it involve all hosts or is it limited to one host?
- Are there special load or configuration conditions under which you encounter a specific problem?
- Does the problem affect a single user? Multiple users? Your LAN?

Troubleshooting Techniques and Tools

1.2 Isolating Problems

The following steps can help you isolate your problem and determine a solution.

1. Check connectivity. (Section 1.2.1)
2. Check network interface parameters. (Section 1.2.2)
3. Check the IP address to Ethernet address translation tables. (Section 1.2.3)
4. Examine network statistics. (Section 1.2.4)
5. Monitor network traffic. (Section 1.2.5)
6. Check name server operation. (Section 1.2.6)
7. Check the route to a remote host. (Section 1.2.7)
8. Check the routes known to a gateway. (Section 1.2.8)
9. Check whether the network services have been enabled. (Section 1.2.9)
10. Look for application errors or interoperability issues.

Table 1–1 summarizes the tools you use to obtain information about network operations. The following sections describe each tool in detail.

Table 1–1 Diagnostic Tools

Diagnostic Tool	Function
arp	Controls and displays ARP tables.
dig	Sends domain name query packets to name servers.
ifconfig	Configures or displays network interface parameters, redefines an address for a particular interface, or sets options such as an alias list, broadcast address, or access filter. Use to detect incorrect IP addresses, subnet masks, and broadcast addresses.
ndc	Allows the name server administrator to send messages to a name server to start, stop, and restart BIND; to dump the BIND database; to check the status of the BIND process; and to change the tracing level.
netstat	Displays network statistics of sockets, data link counters, specified protocols or aliases, network interfaces, and a host's routing table.
nslookup	Provides the ability to directly query a name server and retrieve information. Use NSLOOKUP to determine whether your local name server is running correctly or to retrieve information from remote name servers.
ping	Indicates a host is reachable, and displays statistics about packet loss and delivery time.
route	Allows the user to manipulate the network routing tables manually.
sysconfig	Displays and maintains the various network subsystem attributes.
TCPTRACE	Traces packets going in and out of the system. To run the trace utility, enter the DCL command TCPTRACE.
tracert	Displays the route of an IP packet sent from the local host to a remote host.

To enter a command at the system prompt, execute the `SYS$STARTUP:TCPIP$DEFINE_COMMANDS.COM` file. This file defines each tool as a foreign command.

See Appendix A for complete reference information about these diagnostic tools.

1.2.1 Testing Connectivity Between Network Hosts

Use the `ping` command to test whether you can reach a remote host from your local system. The `ping` command sends an Internet Control Message Protocol (ICMP) echo request to the specified host name or host address. When received by a host, an ICMP reply is returned to the requester.

When using the `ping` command to isolate a problem, you should first test the `localhost` to verify that the system can communicate with itself. For example:

```
TCPIP> ping localhost
PING LOCALHOST (127.0.0.1): 56 data bytes
64 bytes from 127.0.0.1: icmp_seq=0 ttl=64 time=1 ms
64 bytes from 127.0.0.1: icmp_seq=1 ttl=64 time=1 ms
64 bytes from 127.0.0.1: icmp_seq=2 ttl=64 time=0 ms
64 bytes from 127.0.0.1: icmp_seq=3 ttl=64 time=0 ms

----LOCALHOST PING Statistics----
4 packets transmitted, 4 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 0/0/1 ms
TCPIP>
TCPIP> ping 127.0.0.1
PING 127.0.0.1 (127.0.0.1): 56 data bytes
64 bytes from 127.0.0.1: icmp_seq=0 ttl=64 time=1 ms
64 bytes from 127.0.0.1: icmp_seq=1 ttl=64 time=1 ms
64 bytes from 127.0.0.1: icmp_seq=2 ttl=64 time=0 ms
64 bytes from 127.0.0.1: icmp_seq=3 ttl=64 time=0 ms

----127.0.0.1 PING Statistics----
4 packets transmitted, 4 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 0/0/1 ms
```

The output from this `ping` command shows that the system is able to send a message down and then back up the protocol stack through the loopback address. The host address `127.0.0.1` and its associated host name, `localhost`, are the loopback address of the local host. This address was devised so that software could use common code to address local processes as well as remote processes. If the command output shows that it received a message for every message it transmitted, then you can be sure that the network software is up and running and that your system should be able to communicate with remote systems.

If you do not receive output similar to that shown in the example, then one of the following conditions may exist:

- TCP/IP Services may not be running.
- A definition for `localhost` is missing from the local database.

If the `ping` command for `localhost` does not respond correctly, try the `ping` command with the IP address `127.0.0.1`. If this command displays correct output, the TCPIP database is missing a definition for `localhost`.

If `localhost` returns the data correctly at this point, use the `ping` command to test another host on the same local network. If you are able to reach this host, then test remote hosts farther and farther away from the local host.

Troubleshooting Techniques and Tools

1.2 Isolating Problems

If the remote host does not respond to the request, the ping command displays the following message:

```
TCPIP> ping a7ulkt
ping: unknown host a7ulkt
%SYSTEM-F-UNREACHABLE, remote node is not currently reachable
```

If you used an IP address in the ping command, the output may be:

```
TCPIP> ping 10.10.22.1
PING 10.10.22.1 (10.10.22.1): 56 data bytes

----10.10.22.1 PING Statistics----
4 packets transmitted, 0 packets received, 100% packet loss
%SYSTEM-F-TIMEOUT, device timeout
```

These error messages could indicate that:

- There is no host with the specified host name.
- Using the host file or DNS/BIND, the system was not able to resolve the specified host name to an IP address.
- There is no host with that IP address.
- The host is down and not responding.

The following sample shows the ping statistics displayed:

```
TCPIP> ping chester
PING chester (16.20.208.53): 56 data bytes
64 bytes from 16.20.208.53: icmp_seq=0 ttl=64 time=0 ms
64 bytes from 16.20.208.53: icmp_seq=1 ttl=64 time=1 ms
64 bytes from 16.20.208.53: icmp_seq=2 ttl=64 time=0 ms
64 bytes from 16.20.208.53: icmp_seq=3 ttl=64 time=1 ms

----chester PING Statistics----
4 packets transmitted, 4 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 0/0/1 ms
```

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.

If you do not specify command options, the ping command displays the results of each ICMP request in sequence, the number of bytes received from the remote host, and the round-trip time on a per-request basis.

Use the output from the ping command to help determine the cause of direct and indirect routing problems such as host is unreachable, connection timed out, and network is unreachable.

This command helps you decide whether further testing is required and where. For example, if someone reports a problem connecting to a remote host, but ping shows packets traveling to the remote system and back, the problem probably resides in the upper (application) layer protocols (such as FTP, TELNET), or the user introduced the error to the application.

If the packets do not make the round trip, the problem probably resides in the lower layers, and perhaps indicates a misconfigured interface or other configuration or routing problems.

When preliminary testing indicates a problem in the lower layers, the next step is to test the network interfaces and routing. Use the ifconfig, netstat, and arp commands for these purposes (see Appendix A).

1.2.1.1 Using ping on a Multihomed Host

If you suspect one of the interfaces is down, you can test the interface by using:

- The ping command on another system to ping the IP address of the suspect interface. Because TCP/IP Services sends messages via the first matching entry in the routing table, you might not be exercising the interface if you test on the local system.
- The traceroute -s src_addr command on the local system. This command uses the IP address specified with the -s flag in outgoing probe packets. If the command fails, the error message may indicate that the problem is with the interface.

1.2.2 Checking the Network Interface Parameters

Use the ifconfig command to check the configuration of a network interface. A common problem is a misconfigured subnet mask or incorrect IP address. Be sure to check the values of these parameters.

To display configuration information for all interfaces, enter the following command:

```
TCPIP> ifconfig -a
LO0: flags=100c89<UP,LOOPBACK,NOARP,MULTICAST,SIMPLEX,NOCHECKSUM>
      inet 127.0.0.1 netmask ffffffff ipmtu 4096
TN0: flags=80<NOARP>
WE0: flags=c43<UP,BROADCAST,RUNNING,MULTICAST,SIMPLEX>
      HWaddr aa:0:4:0:71:f8
      inet 10.10.2.1 netmask ffffffff broadcast 10.10.2.255 ipmtu 1500
```

For example, to display the configuration for interface WF0, enter the following command:

```
TCPIP> ifconfig WF0
```

The system displays the following information:

```
WF0: flags=c43<UP,BROADCAST,RUNNING,MULTICAST,SIMPLEX>
      inet 16.20.208.100 netmask ffff0000 broadcast 16.20.255.255 ipmtu 4470
      inet6 fe80::200:f8ff:febd:bc22
      inet6 3ffe:1200:4120:1000:200:f8ff:febd:bc22
```

The first line of this display shows the interface characteristics. The interface should be UP and RUNNING (exceptions to this are the LO0 and TN0 interfaces). The pertinent fields in this display show the interface's IP address, the subnet mask, the broadcast mask, and the maximum transmit unit.

You can also obtain information using the following commands:

```
TCPIP> SHOW INTERFACE
```

Interface	IP_Addr	Network mask	Packets		MTU
			Receive	Send	
LO0	127.0.0.1	255.0.0.0	137	137	4096
WE0	10.10.2.1	255.255.255.0	5089	4191	1500

```
TCPIP> SHOW CONFIGURATION INTERFACE
```

```
Interface: LO0
  IP_Addr: 127.0.0.1      NETWRK: 255.0.0.0      BRDCST:
Interface: WE0
  IP_Addr: 10.10.2.1     NETWRK: 255.255.255.0 BRDCST: 10.10.2.255
```

Troubleshooting Techniques and Tools

1.2 Isolating Problems

If you are not familiar with IP addressing and the concepts of subnet and broadcast masks, review the information in *DIGITAL TCP/IP Services for OpenVMS Concepts and Planning* before proceeding with troubleshooting tasks.

1.2.3 Displaying and Modifying the Internet-to-Ethernet Translation Tables

Use the `arp` utility or the `SHOW ARP` management command to check the IP address to Ethernet address translation entries in the Address Resolution Protocol (ARP) table. This is useful if you think incorrect entries are being added to the ARP table. For example, if you enter a command and an unexpected host responds, you may have two systems defined with the same IP address in the ARP table.

To display entries in the ARP table, enter the following command:

```
$ TCPIP SHOW ARP
Cnt  Flags    Timer Host                Phys Addr
 1: UCS      451 160.20.0.10          08-00-2b-39-c7-ac
 2: UC        0 160.20.0.100         aa-00-04-00-8d-13
 3: UC        3 160.20.0.173         00-00-f8-45-a0-b4
 4: UC       14 160.20.32.94         00-00-f8-00-f7-41
 5: UC       50 160.20.64.69         00-d0-b7-19-78-a4
 6: UCS       9 160.20.64.132        00-50-8b-72-7f-ff
 7: UCS     150 160.20.80.124        00-50-8b-4d-91-b3
```

The following TCP/IP Services management commands allow you to configure the hardware addresses for remote IP addresses:

- `SET ARP` allows you to add the hardware address to the ARP table.
- `SET PROTOCOL ARP` allows you specify the time interval for ARP to hold information in its cache. Use the `SET CONFIGURATION PROTOCOL ARP` command to set this information in the permanent configuration database.

For more information about these procedures, refer to the *Compaq TCP/IP Services for OpenVMS Management* manual.

For information about using the `arp` utility, refer to Appendix A.

1.2.4 Examining Network Statistics

Use the `netstat` utility or the `SHOW INTERFACE` command to check interface and protocol statistics, per-connection status, and memory buffer use. Look for bad checksums, excessive retransmissions, dropped packets, out-of-order packets, and lost-carrier errors.

For example:

```
TCPIP> netstat -i
Name Mtu Network Address Ipkts Ierrs Opkts Oerrs Coll
TN0* 1280 Link Link#2 0 0 0 0 0
WF0 4470 Link 00:00:f8:cd:1e:e4 48855499 0 2035244 0 0
WF0 4470 16.20 ucxaxp 48855499 0 2035244 0 0
LO0 4096 Link Link#1 165084 0 165084 0 0
LO0 4096 127 LOCALHOST 165084 0 165084 0 0
```

The netstat command displays information used to diagnose failures. Some problems to look for include:

- **Network problems**
If the netstat -i command shows an excessive number of input errors (Ierrs), output errors (Oerrs), or collisions (Coll), you may be experiencing a network problem, such as improperly connected cables or a saturated Ethernet.
- **Memory problems**
Use netstat -m to see whether the network is using an excessive amount of memory in proportion to the total amount of memory installed on the system. If several memory requests are delayed or denied, this means that your system was temporarily short of physical memory.
- **Network connections**
Use netstat -an to check allocated sockets (each socket results in a network connection).

1.2.5 Monitoring Network Traffic

The trace utility (TCPTRACE) is a tool you can use to trace packets going in and out of the system. To run the trace utility, enter the DCL command TCPTRACE. Use the qualifiers listed in the command reference section to customize tracing for your particular problem. For example:

```
$ TCPTRACE HOST1 /FULL /PORT=REMOTE=21
$ TCPTRACE HOST2 /PORT=(LOCAL=23, REMOTE=1056) /FULL /PACKETS=30 /OUTPUT=TELNET_TRACE.TXT
```

The following sample is a TCPTRACE display:

```
TCPIP INTERNet trace RCV packet seq # = 1 at 23-OCT-1998 15:19:33.29
IP Version = 4, IHL = 5, TOS = 00, Total Length = 217 = ^x00D9
IP Identifier = ^x0065, Flags (0=0,DF=0,MF=0),
  Fragment Offset = 0 = ^x0000, Calculated Offset = 0 = ^x0000
IP TTL = 32 = ^x20, Protocol = 17 = ^x11, Header Checksum = ^x8F6C
IP Source Address = 16.20.168.93
IP Destination Address = 16.20.255.255

UDP Source Port = 138, UDP Destination Port = 138
UDP Header and Datagram Length = 197 = ^x00C5, Checksum = ^x0E77

5DA81410 8F6C1120 00000065 D9000045 0000 E...awe.....l....]
| 0E77C500 8A008A00 | FFFF1410 0010 .....w.
```

For more information about using TCPTRACE, see Appendix A.

Troubleshooting Techniques and Tools

1.2 Isolating Problems

1.2.6 Checking Name Server Operation

After verifying that the underlying transport is working, check to see whether the remote host can be reached by its host name. If your name server resides on a remote system, make sure your resolver configuration specifies that system. To determine whether the resolver is pointing to the correct server, enter the following command:

```
TCPIP> SHOW NAME_SERVICE
BIND Resolver Parameters
Local domain: lkg.dec.com
System
State:      Started, Enabled
Transport:  UDP
Domain:     lkg.dec.com
Retry:      4
Timeout:    4
Servers:    rufus.lkg.dec.com, peach.lkg.dec.com
Path:       lkg.dec.com
Process
State:      Enabled
Transport:
Domain:
Retry:
Timeout:
Servers:
Path:
```

Make sure the remote servers are reachable (using ping) and that they are valid name servers.

If your name server resides on the local system, use the SHOW NAME_SERVICE command to make sure your resolver points to localhost.

Troubleshooting Techniques and Tools

1.2 Isolating Problems

Next, verify that the TCPIP\$BIND process is enabled and running. First, enter the following command to determine whether TCPIP\$BIND is enabled:

```
TCPIP> SHOW SERVICE
```

Service	Port	Proto	Process	Address	State
BIND	53	TCP,UDP	TCPIP\$BIND	0.0.0.0	Enabled
DHCP	67	UDP	TCPIP\$DHCP	0.0.0.0	Enabled
DIOSERVER	1451	TCP	CLM	0.0.0.0	Disabled
ECHO	7	TCP	MULTI	0.0.0.0	Disabled
ESNMP	705	UDP	ESNMP	0.0.0.0	Disabled
FINGER	79	TCP	TCPIP\$FINGER	0.0.0.0	Enabled
FTP	21	TCP	TCPIP\$FTP	0.0.0.0	Enabled
HELLO	12345	TCP	HELLO_WORLD	0.0.0.0	Disabled
JOHN	520	UDP	UCX\$ROUTER	0.0.0.0	Disabled
LBROKER	6570	UDP	TCPIP\$LBROKER	0.0.0.0	Disabled
LPD	515	TCP	TCPIP\$LPD	0.0.0.0	Enabled
MATT	5432	TCP	TCPIP\$RLOGIN	0.0.0.0	Disabled
METRIC	570	UDP	TCPIP\$METRIC	0.0.0.0	Enabled
MOUNT	10	TCP,UDP	TCPIP\$MOUNTD	0.0.0.0	Enabled
NFS	2049	UDP	TCPIP\$NFS	0.0.0.0	Enabled
NOTES	3333	TCP	NOTESRVR	0.0.0.0	Enabled
NTP	123	UDP	TCPIP\$NTP	0.0.0.0	Enabled
PCNFS	5151	TCP,UDP	TCPIP\$PCNFSD	0.0.0.0	Enabled
POP	110	TCP	TCPIP\$POP	0.0.0.0	Enabled
PORTMAPPER	111	TCP,UDP	TCPIP\$PORTM	0.0.0.0	Enabled
REXEC	512	TCP	TCPIP\$REXEC	0.0.0.0	Enabled
RLOGIN	513	TCP	not defined	0.0.0.0	Enabled
RSH	514	TCP	TCPIP\$RSH	0.0.0.0	Enabled
SMTP	25	TCP	TCPIP\$SMTP	0.0.0.0	Enabled
SNMP	161	UDP	TCPIP\$SNMP	0.0.0.0	Enabled
TELNET	23	TCP	not defined	0.0.0.0	Enabled
TFTP	69	UDP	TCPIP\$TFTP	0.0.0.0	Enabled
XDM	177	UDP	TCPIP\$XDM	0.0.0.0	Enabled

If the BIND process is enabled, it will appear in the display.

Then determine whether the BIND process is running by entering the following command:

```
$ SHOW SYSTEM /NETWORK
OpenVMS V7.1-1H2 on node RUFUS 27-JUN-2000 16:45:46.84 Uptime 16 01:55:35
  Pid Process Name State Pri I/O CPU Page flts Pages
2FC0021F TCPIP$NTP LEF 10 2042786 0 00:02:03.43 657 190 N
2FC00221 TCPIP$LBROKER LEF 9 3779921 0 00:06:27.51 652 271 N
2FC05046 TCPIP$POP_1 HIB 10 243688 0 00:00:48.42 955 598 N
2FC00289 TCPIP$PORTM LEF 10 13289 0 00:00:03.23 614 189 N
2FC0628F TCPIP$RE_BG1879 LEF 6 1647 0 00:00:00.96 1709 612 N
2FC0089A NFSS$SERVER LEF 10 89284 0 00:00:19.28 978 580 N
2FC06C9E NOTES$00CD_2* HIB 6 208844 0 00:01:22.65 1932 152 N
2FC03EC7 TCPIP$BIND_1 LEF 10 515297 0 00:01:26.06 972 322 N
2FC01CF6 TCPIP$PCNFSD LEF 10 326 0 00:00:00.27 660 228 N
$
```

If the TCPIP\$BIND_1 process is not running, look for errors in the SYSS\$SPECIFIC:[TCPIP\$BIND]TCPIP\$BIND_RUN.LOG file.

To reduce the possibility of a name server being unavailable, you might configure more than one name server on your network. This way, if the primary name server is unreachable or unresponsive, the resolver can query the other name server.

Troubleshooting Techniques and Tools

1.2 Isolating Problems

1.2.7 Checking the Route to a Remote Host

If you receive “network unreachable” messages, you may be experiencing a routing problem. You can easily detect whether the problem is with your local routing table by doing the following:

- Enter a `netstat -rn` or `SHOW ROUTE` command.

Display the routing table, then compare the output to the routing table of a properly running system. Make sure there is a default route defined and that the IP address listed in the gateway column for the default route and the local host are in the same subnet. The default route specifies the gateway to use when a route is not explicitly defined for the destination IP address.

For example, enter the following command:

```
TCPIP> netstat -rn

Routing tables
Destination      Gateway          Flags    Refs      Use  Interface
-----
Route Tree for Protocol Family 2
default          16.20.0.173     UG       17  1526068  WE0
10.10/16         16.20.208.154  UGS      0   204911  WE0
10.10.39/25      10.10.39.2     U        2   17942   BE0
16.20/16         16.20.208.100  U       45  6219676  WE0
16.20/16         16.20.208.208  U        0     0       WE0
127.0.0.1        127.0.0.1      UH       1   69844   LO0

Route Tree for Protocol Family 26
::1              Link#1          UH       0     0       LO0
ff01::/16        Link#1          U        0     0       LO0
```

- To display a default route using the TCP/IP Services management commands, enter one of the following commands:

```
$ TCPIP SHOW ROUTE /PERMANENT /DEFAULT
$ TCPIP SHOW ROUTE /DEFAULT
```

The following example shows typical output from these two commands:

```
$ TCPIP SHOW ROUTE /PERMANENT /DEFAULT
                                PERMANENT
Type          Destination          Gateway
PN    0.0.0.0          rufus.lkg.dec.com

$ TCPIP SHOW ROUTE /DEFAULT
                                DYNAMIC
Type          Destination          Gateway
DN    0.0.0.0          10.10.2.66
$
```

To set a default route, enter a command similar to the following:

```
$ TCPIP SET ROUTE /DEFAULT /GATE=n.n.n.n
```

You can also set a default route by running the `TCPIP$CONFIG` procedure and selecting option 1 for Core, and then option 3 for Routing. `TCPIP$CONFIG` prompts with:

```
* Do you want to configure dynamic ROUTED or GATED routing [NO]:
```

Take the default value by pressing the Enter key. TCPIP\$CONFIG then displays the current configuration and asks whether you want to reconfigure a default route:

The current configuration for the default route is:

```
PERMANENT

Type          Destination          Gateway
PN    0.0.0.0          rufus.lkg.dec.com

* Do you want to reconfigure a default route [YES]:
Enter the Default Gateway host name []:
```

- Next, use ping to see whether you can reach the routing gateway.

1.2.8 Checking the Routes Known to a Gateway

The `tracert` command helps you locate problems between the local host and the remote destination by tracing the route of UDP packets from the local host to a remote host. Tracing attempts to determine the name and IP address of each gateway along the route to the remote host.

The `tracert` command works by sending UDP packets with small time-to-live (TTL) values and an invalid port number to the remote system. The TTL values increase in increments of one for each group of three UDP packets sent. When a gateway receives a packet, it decrements the TTL. If the TTL is zero, the packet is not forwarded, and an ICMP “time exceeded” message is returned.

Intermediate gateways are detected when they return an ICMP “time exceeded” message. When `tracert` receives an “invalid port” message, it knows that it reached the remote destination. (`tracert` operates by intentionally using an invalid port.) When `tracert` receives this message, it knows it has reached the destination host and terminates the trace. In this way, `tracert` develops a list of gateways starting at one hop away, and increasing one hop at a time until the remote host is reached.

For more information about using `tracert`, see Appendix A.

1.2.9 Determine Whether Network Services Are Available

The auxiliary server functions like the UNIX internet daemon (`inetd`) by managing access to the network services. The auxiliary server assigns standard port numbers to services such as the BOOTP, SMTP, or FTP servers, and starts the appropriate image after receiving an incoming request.

To verify correct operation of a service, you need to verify that the service:

- Has an entry in the service database
- Has the correct attributes defined
- Account has the correct privileges
- Is enabled
- Is started

Troubleshooting Techniques and Tools

1.2 Isolating Problems

1.2.9.1 Displaying the Service Database

To display the services database, enter the `SHOW SERVICE` command. For example:

```
TCPIP> SHOW SERVICE
```

1	2	3	4	5	6
Service	Port	Proto	Process	Address	State
FINGER	79	TCP	TCPIP\$FINGER	0.0.0.0	Disabled
FTP	21	TCP	TCPIP\$FTP	0.0.0.0	Enabled
LPD	515	TCP	TCPIP\$LPD	0.0.0.0	Enabled
MOUNT	10	UDP	TCPIP\$NFS_M	0.0.0.0	Enabled
NFS	2049	UDP	TCPIP\$NFS	0.0.0.0	Enabled
NTP	123	UDP	TCPIP\$NTP	0.0.0.0	Enabled
PCNFS	5151	TCP,UDP	TCPIP\$PCNFSD	0.0.0.0	Enabled
POP	110	TCP	TCPIP\$POP	0.0.0.0	Enabled
PORTMAPPER	111	TCP,UDP	TCPIP\$PORTM	0.0.0.0	Enabled
REXEC	512	TCP	TCPIP\$REXEC	0.0.0.0	Enabled
RLOGIN	513	TCP	not defined	0.0.0.0	Enabled
RSH	514	TCP	TCPIP\$RSH	0.0.0.0	Enabled
SMTP	25	TCP	TCPIP\$SMTP	0.0.0.0	Enabled
SNMP	161	UDP	TCPIP\$SNMP	0.0.0.0	Enabled
TELNET	23	TCP	not defined	0.0.0.0	Enabled
TFTP	69	UDP	TCPIP\$TFTP	0.0.0.0	Enabled

- 1 This column lists those services with entries in the TCPIP services database. If not listed in this column, the service was never enabled during the configuration procedure (using `TCPIP$CONFIG.COM`). To enable additional services, run the `TCPIP$CONFIG` procedure.
- 2 This column lists the port on which the service listens for connection requests. The port number is either the well-known port number for the service or an ephemeral port number assigned when the socket is assigned a protocol address.
- 3 This column lists the TCP/IP protocol that the service uses to communicate with the client process.
- 4 This column lists the process name for the service. If you use the `DCL` command `SHOW SYSTEM /NETWORK`, this is the process name you should see if the process is running.
- 5 This column lists the IP address of the interface on which the service accepts connection requests. IP address 0.0.0.0 indicates that the service will accept connection requests received on any of the local interfaces.
- 6 This column lists whether the service is enabled or disabled. The term enabled indicates that the next time TCP/IP Services starts, TCP/IP Services starts all services that are marked in the service database as enabled. In this example, of the services listed, all services except `finger` will start the next time TCP/IP Services restarts.

Note

In this example, the `finger` service was configured with `TCPIP$CONFIG`. However, at some point, `finger` was disabled either by a TCPIP management command or by an incremental shutdown of the service.

1.2.9.2 Displaying Service Attributes

Each service should have the following items defined in the services database:

- OpenVMS user account (also in user authorization file [UAF])
- Unique port number
- Protocol
- Name and location of the startup command procedure and log file
- Service parameters (for example, timeouts, privileges)
- Flags

If these items are not defined correctly, or if the service account privileges and file protections are not assigned correctly, the service will fail to respond to an incoming request. This failure may be logged in the service-specific log file.

To display information about a service, enter the TCPIP command `SHOW SERVICE /FULL` and specify the service name. For example:

```
$ TCPIP
TCPIP> SHOW SERVICE /FULL TELNET

Service: TELNET                                     1
State:      Enabled
Port:       23      Protocol:  TCP      Address:  0.0.0.0
Inactivity: 1      User_name: not defined Process:  not defined
Limit:      57     Active:    12      Peak:     14

File:       not defined
Flags:      Listen Rtty

Socket Opts: Keepalive Rcheck Scheck                2
Receive:    3000      Send:          3000

Log Opts:   Actv Dactv Conn Error Logi Logo Mdfy Rjct 3
File:       not defined

Security                                         4
Reject msg: not defined
Accept host: 0.0.0.0
Accept netw: 0.0.0.0
TCPIP>
```

- 1 This section displays information about the service: service name, process name, user name, port and interface on which the service is listening, whether the service is enabled or disabled, and the number of copies of the service that can run at one time.
- 2 This section displays the socket options that the service uses. The service's socket options can be changed dynamically, though it is unlikely that someone would change them. If you suspect that improper socket options are in effect, you can reestablish the default values by disabling the service, running `TCPIP$CONFIG`, and then enabling the service.
- 3 This section displays the name of the log file that receives event messages and the events that the service will log. Checking the log file may indicate the cause of a problem.
- 4 This security section displays a list of hosts and networks that are specifically given or denied access to the service. If one system is unable to access a service, check this section to see whether the system or its associated network is being denied the service.

Troubleshooting Techniques and Tools

1.2 Isolating Problems

1.2.9.3 Verifying Process Privileges

To check the privileges associated with a service's process, enter a command for the process, as follows:

```
$ INSTALL LIST/FULL TCPIP$SMTP_RECEIVER
DISK$VMS721:<SYS0.SYSCOMMON.SYSEXE>.EXE
  TCPIP$SMTP_RECEIVER;1
      Open Hdr Shared  Prv
Entry access count      = 20
Current / Maximum shared = 1 / 1
Global section count    = 1
Privileges = SYSPRV
Authorized = SYSPRV

$ INSTALL LIST/FULL TCPIP$FTP_CHILD
DISK$VMS721:<SYS0.SYSCOMMON.SYSEXE>.EXE
  TCPIP$FTP_CHILD;1
      Open Hdr Shared  Prv
Entry access count      = 42
Current / Maximum shared = 1 / 3
Global section count    = 1
Privileges = PSWAPM OPER
Authorized = PSWAPM OPER
```

1.2.9.4 Verifying Account Privileges

To determine the privileges associated with the service's account, run the OpenVMS Authorize utility and then use the SHOW command with the process name of the service, as follows:

```
A72KT: SET DEFAULT SYS$SYSTEM
A72KT: RUN AUTHORIZE
UAF> SHOW TCPIP$SNMP

Username: TCPIP$SNMP          Owner: TCPIP$SNMP
Account: TCPIP                UIC: [3655,13] ([TCPIP$AUX,TCPIP$
NMP])
CLI: DCL                      Tables: DCLTABLES
Default: SYS$SYSDEVICE:[TCPIP$SNMP]
LGICMD: LOGIN
Flags: Restricted
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: 0
Pwdlifetime: 90 00:00 Pwdchange: (pre-expired)
Last Login: (none) (interactive), 7-AUG-2000 12:45 (non-interactive)
Maxjobs: 0 Fillm: 50 Byt1m: 52200
Maxacctjobs: 0 Shrfillm: 0 Pbyt1m: 0
Maxdetach: 0 B101m: 18 JTquota: 4096
Prclm: 8 D101m: 18 WSdef: 350
Prio: 8 AST1m: 100 WSquo: 512
Queprio: 4 TQElm: 15 WSextent: 512
CPU: (none) Enqlm: 100 Pgflquo: 10240
Authorized Privileges:
  NETMBX TMPMBX
Default Privileges:
  NETMBX TMPMBX
```

1.2.9.5 Looking for OPCOM Messages

The following is another method of detecting failure of the auxiliary server to start a service:

1. Enter the following commands:

```
$ SET PROCESS /PRIVILEGE=OPER      1
$ SET AUDIT /ALARM /ENABLE=FAILURE  2
$ REPLY/ENABLE=NETWORK             3
```

- 1 Enables your process with operator privileges.
 - 2 Enables security auditing to log unsuccessful object attempts.
 - 3 Establishes your terminal as an operator's terminal.
2. After you enter these commands, have a remote user try to connect to the service on the local system.
 3. Review any OPCOM messages for indications of the problem.

Tuning Techniques

This chapter describes how to adjust TCP/IP variables to improve network performance.

2.1 Kernel Subsystems

The TCP/IP Services kernel contains the following subsystems:

- inet
- net
- socket
- iptunnel
- ipv6
- snmpinfo

Each subsystem has attributes that you can change to affect the performance of the network. You can display and modify these attributes values by using the `sysconfig` command.

The following sections describe how to perform these tasks:

- Display kernel subsystems and attributes (Section 2.1.1)
- Test attribute support (Section 2.1.2)
- Modify attribute values (Section 2.1.4)

2.1.1 Displaying Subsystems and Attributes

To display kernel subsystems, enter the following command:

```
$ TCPIP
TCPIP> sysconfig -s
inet: loaded and configured
net: loaded and configured
socket: loaded and configured
iptunnel: loaded and configured
ipv6: loaded and configured
snmpinfo: loaded and configured
TCPIP>
```

To display the attributes of a particular subsystem, enter a command similar to the following:

Tuning Techniques

2.1 Kernel Subsystems

```
TCPIP> sysconfig -q socket
socket:
sbcompress_threshold = 0
sobacklog_drops = 0
sobacklog_hiwat = 3
somaxconn = 1024
somaxconn_drops = 0
sominconn = 0
TCPIP>
```

2.1.2 Testing Attribute Support

To determine support for an attribute, use the `sysconfig -q subsystem [attribute]` command.

If you do not specify an attribute, the system displays all the subsystem attributes that can be modified with the `sysconfig` command. If the subsystem is not configured, `sysconfig` displays a message similar to the following:

```
framework error: subsystem 'inet' not found
```

If you specify an attribute, the `sysconfig` command displays only information about that attribute. For example:

```
# sysconfig -q inet tcbhashsize
inet:
tcbhashsize = 32
```

If the attribute is not supported or if it cannot be accessed by using `sysconfig`, `sysconfig` displays a message similar to the following message:

```
inet:
tcbhashsize = unknown attribute
```

For more information about using the `sysconfig` command, see Appendix A.

2.1.3 Displaying Attribute Values

Use the following methods to display attribute values:

- The `sysconfig -q subsystem [attribute]` command displays the permanent value for attributes. If you do not specify an attribute, the system displays all the subsystem attributes that can be modified with the `sysconfig` command.

If you specify an attribute, information about only that attribute is displayed. For example:

```
TCPIP> sysconfig -q inet tcp_rexmtmax
inet:
tcp_rexmtmax = 128
TCPIP>
```

- The `sysconfig -Q subsystem [attribute]` command displays the maximum and minimum values that you can specify for attributes. If you do not specify an attribute, the system displays all the subsystem attributes that can be modified with the `sysconfig` command.

If you specify an attribute, information about only that attribute is displayed. For example:

```
TCPIP> sysconfig -Q inet tcp_rexmtmax
inet:
tcp_rexmtmax - type=INT op=CRQ min_val=1 max_val=2147483647
TCPIP>
```

For more information about the `sysconfig` command, see the Appendix A .

2.1.4 Modifying Attribute Values

The `/etc/sysconfigtab` subsystem attribute database file contains modifications to the default attribute values. Various methods are available to modify attribute values.

Note

Use the `sysconfig -r` command to modify attribute values in the `sysconfigtab` file. Do not modify the file manually.

2.1.5 Modifying Attributes Temporarily

You may be able to modify an attribute temporarily by changing only its current value. This allows you to determine whether modifying an attribute will improve your system performance. Not all attributes can be changed dynamically.

Temporary modifications are lost when you reboot the system.

To modify an attribute's current value, use the following method:

- The `sysconfig -r` command allows you to modify the current value of an attribute if the attribute supports this operation:

```
sysconfig -r subsystem attribute=value
```

For example:

```
TCPIP> sysconfig -r inet tcp_keepinit=30
tcp_keepinit: reconfigured
TCPIP>
```

For information about the `sysconfig` command, see Appendix A.

2.2 Modifying Kernel Subsystems

Most resources used by the network subsystem are allocated and adjusted dynamically. However, you can make some adjustments to improve performance.

Table 2–1 summarizes the adjustments you can make, lists performance benefits and the adjustments that will achieve them, along with the tradeoffs (where applicable) associated with each adjustment.

Table 2–1 Network Tuning Guidelines

Performance Benefit	Tuning Adjustment	Tradeoff
Reduce the number of dropped incoming connection requests.	Increase the maximum number of pending TCP connections (Section 2.2.1.1).	Consumes memory resources.
Allow each server socket to handle more SYN packets simultaneously.	Increase the minimum number of pending TCP connections (Section 2.2.1.2).	Consumes memory resources.
Allow for a larger socket buffer.	Increase the maximum socket buffer size (Section 2.2.1.3).	Consumes memory. If you have a large number of sockets, memory consumption could be of concern.

(continued on next page)

Tuning Techniques

2.2 Modifying Kernel Subsystems

Table 2–1 (Cont.) Network Tuning Guidelines

Performance Benefit	Tuning Adjustment	Tradeoff
Improve the TCP control block lookup rate and increase the raw connection rate.	Increase the size of the hash table that the kernel uses to look up TCP control blocks (Section 2.2.2.1).	Slightly increases the amount of pooled memory.
Reduce hash table lock contention for SMP systems.	Increase the number of TCP hash tables (Section 2.2.2.2).	Slightly increases the amount of pooled memory.
Improve performance on systems that use large numbers of interface alias.	Increase the size of the kernel interface alias table (Section 2.2.2.3).	None.
Allow partial connections to time out sooner, preventing the socket listen queue from filling up with SYN packets.	Increase the TCP partial connection timeout rate (Section 2.2.2.4).	Setting the <code>tcp_keepinit</code> value too low can cause connections to be broken prematurely.
Prevent premature retransmissions and decrease congestion.	Reduce the TCP retransmission rate (Section 2.2.2.5).	A long retransmit time is not appropriate for all configurations.
Clean up sockets that do not exit cleanly when the <code>keepalive</code> interval expires.	Enable TCP <code>keepalive</code> functionality (Section 2.2.2.6).	None.
Free connection resources sooner.	Make the TCP connection context time out more quickly at the end of the connection (Section 2.2.2.7).	Reducing the timeout limit increases the potential for data corruption; use caution if you make this adjustment.
Provide TCP and UDP applications with a specific range of ports.	Modify the range of outgoing connection ports (Section 2.2.2.8).	None.
Improve the efficiency of servers that handle remote traffic from many clients.	Disable the use of a PMTU (Section 2.2.2.9).	May reduce server efficiency for LAN traffic.
Allow large socket buffer sizes.	Increase the maximum size of a socket buffer (Section 2.2.1.3).	Consumes memory resources.

The following sections describe in detail how to modify socket subsystem attributes and internet subsystem attributes.

2.2.1 Modifying Socket Subsystem Attributes

The `socket` subsystem attributes control the maximum number of pending connection attempts per server socket (that is, the maximum depth of the listen or SYN queue) and other behavior. You may be able to improve server performance by modifying the `socket` subsystem attributes described in Table 2–2.

Table 2–2 `socket` Subsystem Attributes

Attribute	Description
<code>somaxconn</code>	Controls the maximum number of pending TCP connections.
<code>sominconn</code>	Controls the minimum number of pending TCP connections.
<code>sb_max</code>	Controls the maximum size of a socket buffer.

In addition, the `socket` subsystem attributes `sobacklog_hiwat`, `sobacklog_drops`, and `somaxconn_drops` track events related to socket listen queues. By monitoring these attributes, you can determine whether the queues are overflowing.

2.2.1.1 Increasing the Maximum Number of Pending TCP Connections

The `socket` subsystem attribute `somaxconn` specifies the maximum number of pending TCP connections (the `socket` listen queue limit) for each server socket (for example, for the HTTP server socket). Busy servers often experience large numbers of pending connections. If the listen queue connection limit is too small, incoming connection requests may be dropped. Pending TCP connections can be caused by lost packets in the internet or denial of service attacks.

The default value for `somaxconn` is 1024.

Compaq recommends increasing the `somaxconn` attribute to the maximum value, except on low-memory systems. The maximum value is 65535. Specifying a value that is higher than the maximum value can cause unpredictable behavior.

2.2.1.2 Increasing the Minimum Number of Pending TCP Connections

The `socket` subsystem attribute `sominconn` specifies the minimum number of pending TCP connections (backlog) for each server socket. This attribute controls how many SYN packets can be handled simultaneously before additional requests are discarded. Network performance can degrade if a client saturates a `socket` listen queue with erroneous TCP SYN packets, effectively blocking other users from the queue.

The value of the `sominconn` attribute overrides the application-specific backlog value, which may be set too low for some server software. If you do not have your application source code, you can use the `sominconn` attribute to set a sufficient pending-connection quota.

The default value is 0.

Compaq recommends increasing the value of the `sominconn` attribute to the maximum value of 65535. The value of the `sominconn` attribute should be the same as the value of the `somaxconn` attribute (see Section 2.2.1.1).

2.2.1.3 Increasing the Maximum Size of a Socket Buffer

The `socket` subsystem attribute `sb_max` specifies the maximum size of a socket buffer.

Performance Benefits and Tradeoffs

Increasing the maximum size of a socket buffer may improve performance if your applications can benefit from a large buffer size.

You can modify the `sb_max` attribute without rebooting the system.

When to Tune

If you require a large socket buffer, increase the maximum socket buffer size.

Recommended Values

The default value of the `sb_max` attribute is 128 KB. Increase this value before you increase the size of the transmit and receive socket buffers (see Section 10.2.16).

Tuning Techniques

2.2 Modifying Kernel Subsystems

2.2.2 Modifying Internet Subsystem Attributes

You may be able to improve `inet` subsystem performance by modifying the attributes described in Table 2–3.

Table 2–3 `inet` Subsystem Attributes

Attribute	Description
<code>tcbhashsize</code>	Controls the size of a TCP hash table.
<code>tcbhashnum</code>	Specifies the number of TCP hash tables.
<code>inifaddr_hsize</code>	Controls the size of the kernel interface alias table.
<code>tcp_keepinit</code>	Specifies the TCP partial connection timeout rate.
<code>tcp_rexmit_interval_min</code>	Specifies the rate of TCP retransmissions.
<code>tcp_keepalive_default</code>	Enables or disables the TCP keepalive function.
<code>tcp_msl</code>	Specifies the TCP connection context timeout rate.
<code>ipport_userreserved</code>	Specifies the maximum value for the range of outgoing connection ports.
<code>ipport_userreserved_min</code>	Specifies the minimum value for the range of outgoing connection ports.
<code>pmtu_enabled</code>	Enables or disables use of the PMTU protocol.
<code>ipqs</code>	Specifies the number of IP input queues.
<code>ipqmaxlen</code>	Prevents dropped input packets.

2.2.2.1 Increasing the Size of a TCP Hash Table

You can modify the size of the hash table that the kernel uses to look up Transmission Control Protocol (TCP) control blocks. The `inet` subsystem attribute `tcbhashsize` specifies the number of hash buckets in the kernel TCP connection table (the number of buckets in the `inpcb` hash table).

Performance Benefits and Tradeoffs

The kernel must look up the connection block for every TCP packet it receives, so increasing the size of the table can speed the search and improve performance. This results in a small increase in pooled memory.

You can modify the `tcbhashsize` attribute without rebooting the system.

When to Tune

Increase the number of hash buckets in the kernel TCP connection table if you have an server.

Recommended Values

The default value of the `tcbhashsize` attribute is 512. For servers, set the `tcbhashsize` attribute to 16384.

2.2.2.2 Increasing the Number of TCP Hash Tables

You can increase the number of hash tables the kernel uses to look up TCP control blocks. Because the kernel must look up the connection block for every Transmission Control Protocol (TCP) packet it receives, a bottleneck may occur at the TCP hash table in SMP systems. Increasing the number of tables distributes the load and may improve performance. The `inet` subsystem attribute `tcbhashnum` specifies the number of TCP hash tables.

Performance Benefits and Tradeoffs

For SMP systems, you may be able to reduce hash table lock contention by increasing the number of hash tables that the kernel uses to look up TCP control blocks. This will slightly increase pooled memory.

You cannot modify the `tcblhashnum` attribute without rebooting the system.

When to Tune

Increase the number of TCP hash tables if you have an SMP system that is an server.

Recommended Values

The minimum value of the `tcblhashnum` attribute is 1 (the default). The maximum value is 64. For busy server SMP systems, you can increase the value of the `tcblhashnum` attribute to 16. If you increase this attribute, you should also increase the size of the hash table by a similar factor. See Section 2.2.2.1 for more information.

Compaq recommends that you make the value of the `tcblhashnum` attribute the same as the value of the `inet` subsystem attribute `ipqs`.

2.2.2.3 Increasing the Size of the Kernel Interface Alias Table

The `inet` subsystem attribute `inifaddr_hsize` specifies the number of hash buckets in the kernel interface alias table (`in_ifaddr`).

If a system is used as a server for many different server domain names, each of which is bound to a unique IP address, the code that matches arriving packets to the right server address uses the hash table to speed lookup operations for the IP addresses.

Performance Benefits and Tradeoffs

Increasing the number of hash buckets in the table can improve performance on systems that use large numbers of aliases.

When to Tune

Increase the number of hash buckets in the kernel interface alias table if your system uses large numbers of aliases.

You can modify the `inifaddr_hsize` attribute without rebooting the system.

Recommended Values

The default value of the `inet` subsystem attribute `inifaddr_hsize` is 32; the maximum value is 512.

For the best performance, the value of the `inifaddr_hsize` attribute is always rounded down to the nearest power of 2. If you are using more than 500 interface IP aliases, specify the maximum value of 512. If you are using fewer than 250 aliases, use the default value of 32. For a number of aliases between 250 and 500, use a value that is a power of 2 between 32 and 512.

2.2.2.4 Increasing the TCP Partial Connection Timeout Rate

If increasing the `somaxconn` limit does not prevent the listen queue from filling, or if the default grows to an excessive length, you can make partial connections time out sooner by decreasing the value of the `inet` subsystem attribute `tcp_keepinit`.

The `tcp_keepinit` attribute specifies the amount of time that a partial connection remains on the socket listen queue before it times out.

Tuning Techniques

2.2 Modifying Kernel Subsystems

Performance Benefits and Tradeoffs

Network performance can degrade if a client overfills a socket listen queue with TCP SYN packets, thereby blocking other users from the queue. To eliminate this problem, increase the value of the `sominconn` attribute to its maximum value. If the system continues to drop SYN packets, decrease the value of the `tcp_keepinit` attribute to 30 (15 seconds). Monitor the values of the `sobacklog_drops` and `somaxconn_drops` attributes to determine whether the system is dropping packets. (See Section 2.3.2 for more information about event counters.)

You can modify the `tcp_keepinit` attribute without rebooting the system.

When to Tune

Modify the TCP partial-connection timeout limit only if the value of the `somaxconn_drops` attribute increases often. If this occurs, decrease the value of the `tcp_keepinit` attribute.

Recommended Values

The value of the `tcp_keepinit` attribute is in units of 0.5 seconds. The default value is 150 units (75 seconds). If the value of the `sominconn` attribute is 65535, use the default value of the `tcp_keepinit` attribute.

If you set the value of the `tcp_keepinit` attribute too low, you may prematurely break connections associated with clients on network paths that are slow or network paths that lose many packets. Do not set the value to less than 20 units (10 seconds).

2.2.2.5 Slowing TCP Retransmission Rate

The `inet` subsystem attribute `tcp_rexmit_interval_min` specifies the minimum amount of time before the first TCP retransmission.

Performance Benefits and Tradeoffs

You can increase the value of the `tcp_rexmit_interval_min` attribute to slow the rate of TCP retransmissions, which decreases congestion and improves performance.

You can modify the `tcp_rexmit_interval_min` attribute without rebooting the system.

When to Tune

Not every connection needs a long retransmission time. Usually, the default value is adequate. However, for some wide area networks (WANs), the default retransmission interval may be too small, causing premature retransmission timeouts. This may lead to duplicate transmission of packets and the erroneous invocation of the TCP congestion-control algorithms.

To check for retransmissions, use the `netstat -p tcp` command and examine the output for data packets retransmitted.

Recommended Values

The `tcp_rexmit_interval_min` attribute is specified in units of 0.5 second. The default value is 2 units (1 second).

Do not specify a value that is less than 1 unit. Do not change the attribute unless you fully understand TCP algorithms and your network topology.

2.2.2.6 Enabling the TCP Keepalive Function

The keepalive function enables the periodic transmission of messages on a connected socket in order to keep connections active. Sockets that do not exit cleanly are cleaned up when the keepalive interval expires. If keepalive is not enabled, those sockets continue to exist until you reboot the system.

Applications enable keepalive for sockets by setting the `setsockopt` function's `SO_KEEPALIVE` option. To override programs that do not set keepalive, or if you do not have access to the application sources, use the `inet` subsystem attribute `tcp_keepalive_default` to enable keepalive functionality.

Performance Benefit

Keepalive functionality cleans up sockets that do not exit cleanly when the keepalive interval expires.

You can modify the `tcp_keepalive_default` attribute without rebooting the system. However, sockets that already exist will continue to use old behavior, until the applications are restarted.

When to Tune

Enable keepalive if you require this functionality, and you do not have access to the source code.

Recommended Values

To override programs that do not set keepalive, or if you do not have access to application source code, set the `inet` subsystem attribute `tcp_keepalive_default` to 1 in order to enable keepalive for all sockets.

If you enable keepalive, you can also configure the following TCP options for each socket:

Table 2–4 TCP Keepalive Options

Option	Description
<code>tcp_keepidle</code>	Specifies the amount of idle time, in 0.5-second units, before sending a keepalive probe. The default interval is 75 seconds.
<code>tcp_keepintvl</code>	Specifies the amount of time, in 0.5-second units, between retransmission of keepalive probes. The default interval is 75 seconds.
<code>tcp_keepcnt</code>	Specifies the maximum number of keepalive probes that are sent before the connection is dropped. The default is 8 probes.
<code>tcp_keepinit</code>	Specifies the maximum amount of time, in 0.5-second units, before an initial connection attempt times out. The default is 75 seconds.

2.2.2.7 Increasing the Timeout Rate for TCP Connection Context

The TCP protocol includes a concept known as the Maximum Segment Lifetime (MSL). When a TCP connection enters the `TIME_WAIT` state, it must remain in this state for twice the value of the MSL; otherwise, undetected data errors on future connections can occur. The `inet` subsystem attribute `tcp_msl` determines the maximum lifetime of a TCP segment and the timeout value for the `TIME_WAIT` state.

In some situations, the default timeout value for the `TIME_WAIT` state (60 seconds) is too large, thereby reducing the value of the `tcp_msl` attribute frees connection resources sooner than the default setting.

Tuning Techniques

2.2 Modifying Kernel Subsystems

Performance Benefits and Tradeoffs

You can decrease the value of the `tcp_msl` attribute to make the TCP connection context time out more quickly at the end of a connection. However, this will increase the chance of data corruption.

You can modify the `tcp_msl` attribute without rebooting the system.

When to Tune

Usually, you do not have to modify the timeout limit for the TCP connection context.

Recommended Values

The value of the `tcp_msl` attribute is set in units of 0.5 second. The default value is 60 units (30 seconds), which means that the TCP connection remains in `TIME_WAIT` state for 60 seconds, or twice the value of the MSL.

Do not reduce the value of the `tcp_msl` attribute unless you fully understand the design and behavior of your network and the TCP protocol. It is strongly recommended that you use the default value; otherwise, there is the potential for data corruption.

2.2.2.8 Modifying the Range of Outgoing Connection Ports

When a TCP or UDP application creates an outgoing connection, the kernel dynamically allocates a nonreserved port number for each connection. The kernel selects the port number from a range of values between the value of the `inet` subsystem attribute `ipport_userreserved_min` and the value of the `ipport_userreserved` attribute. Using the default values for these attributes, the range of outgoing ports starts at 49152 and stops at 65535.

Performance Benefits and Tradeoffs

Modifying the range of outgoing connections provides TCP and UDP applications with a specific range of ports.

You can modify the `ipport_userreserved_min` and `ipport_userreserved` attributes without rebooting the system.

When to Tune

If your system requires outgoing ports from a particular range, you can modify the values of the `ipport_userreserved_min` and `ipport_userreserved` attributes.

Recommended Values

The default value of the `ipport_userreserved_min` attribute is 49152. The default value of the `ipport_userreserved` is 65535. The maximum value of each attribute is 65535.

Do not reduce the `ipport_userreserved` attribute to a value that is less than 65535, and do not reduce the `ipport_userreserved_min` attribute to a value that is less than 49152.

2.2.2.9 Disabling Use of the PMTU Protocol

Packets transmitted between servers are fragmented into units of a specific size in order to ease transmission of the data over routers and small-packet networks, such as Ethernet networks. When the `inet` subsystem attribute `pmtu_enabled` is enabled (set to 1, which is the default behavior), the system determines the largest common path maximum transmission unit (PMTU) value between servers and uses it as the unit size. The system also creates a routing table entry for each client network that attempts to connect to the server.

Performance Benefits and Tradeoffs

If a server handles traffic among many remote clients, disabling the use of a PMTU can decrease the size of the kernel routing table, which improves server efficiency. However, on a server that handles local traffic and some remote traffic, disabling the use of a PMTU can degrade bandwidth.

When to Tune

If an server has poor performance and the routing table increases to more than 1000 entries, you should disable the use of PMTU. This is also recommended if you have a server that handles traffic among many remote clients.

Recommended Values

To disable the use of PMTU protocol, set the value of the `pmtu_enabled` attribute to 0.

2.3 Monitoring Servers

Two ways to monitor the behavior of your server and to diagnose performance problems are:

- Using the `netstat` command
This command displays a list of active sockets for each network protocol, information about network routes, and cumulative statistics for network interfaces, including the number of incoming and outgoing packets and the number of packet collisions. The `netstat` command also displays information about memory used for network operations. For more information about monitoring the network, see Section 2.3.1 .
- Monitoring socket subsystem event counters
The socket subsystem attributes `sobacklog_hiwat`, `sobacklog_drops`, and `somaxconn_drops` track events related to socket listen queues. By monitoring these attributes, you can determine if the queues are overflowing. See Section 2.3.2 for information more information about event counters.

2.3.1 Displaying Network Statistics

The `netstat` command displays network statistics, including information about network routes and active sockets for each protocol. The command also displays cumulative statistics for network interfaces, including the number of incoming and outgoing packets and packet collisions, information about memory used for network operations, and statistics related to IP, ICMP, TCP, and UDP protocol layers. You can use the `netstat` command to identify problems by looking for large numbers of bad checksums, retransmissions, and error packets.

Some problems to look for are:

- If the output of the `netstat -i` command shows excessive amounts of input errors (Ierrs), output errors (Oerrs), or collisions (Coll), this may indicate a network problem; for example, cables may not be connected properly or the Ethernet may be saturated.
- The `netstat -m` command shows statistics for network-related data structures. Use this command to determine if the network is using an excessive amount of memory in proportion to the number of sockets in use.

Tuning Techniques

2.3 Monitoring Servers

The following example shows the output of the `netstat -m` command:

```
TCPIP> netstat -m
134 mbufs in use:
    1 mbufs allocated to data
    1 mbufs allocated to packet headers
    24 mbufs allocated to socket structures          1
    35 mbufs allocated to protocol control blocks  2
    18 mbufs allocated to routing table entries
    2 mbufs allocated to socket names and addresses
    13 mbufs allocated to interface addresses
    1 mbufs allocated to OpenVMS kernel table
    2 mbufs allocated to OpenVMS ip multicast address
    2 mbufs allocated to OpenVMS interface multicast address
    1 mbufs allocated to OpenVMS IFNET structure
    3 mbufs allocated to OpenVMS TCPIP Timer structure
    3 mbufs allocated to OpenVMS LAN VCI VCIB structure
    1 mbufs allocated to OpenVMS LAN MCAST_REQ structure
    5 mbufs allocated to OpenVMS SELECT structure
    1 mbufs allocated to OpenVMS ACP AQB
    1 mbufs allocated to OpenVMS ACP INETCB
    1 mbufs allocated to OpenVMS Driver requested REQCB
    19 mbufs allocated to OpenVMS ACP allocated SERV Structure
```

- 1 This line indicates there are 24 sockets in use (1 mbuf allocated for each socket).
- 2 There are two protocol control blocks allocated for each TCP socket and one protocol control block for each UDP socket. The 35 mbuf listed is a mix of PCBs allocated for TCP and UDP sockets. Output from the `TCPIP SHOW DEVICE_SOCKET` command will tell you how many TCP and UDP sockets are allocated.

By comparing the information output from the `netstat -m` and the `TCPIP SHOW DEVICE_SOCKET` command, you can estimate whether the system is using an excessive amount of memory for the number of allocated sockets.

If you sense that TCP/IP Services is using an excessive amount of memory for the number of sockets, there may be a memory leak. Capture the output from the `netstat -m` and the `TCPIP SHOW DEVICE_SOCKET` commands and save for documenting the condition.

The following table shows variations of the `netstat` command that can reveal network problems:

Command	Purpose
<code>netstat -p ip</code>	Checks for bad checksums, length problems, excessive redirects, and packets lost because of resource problems.
<code>netstat -p tcp</code>	Checks for retransmissions, out of order packets, and bad checksums.
<code>netstat -p udp</code>	Looks for bad checksums and full sockets.
<code>netstat -rs</code>	Obtains routing statistics.
<code>netstat -s</code>	Simultaneously displays statistics related to the IP, ICMP, TCP, and UDP protocol layers.
<code>netstat -is</code>	Checks for network device driver errors.

For more information about `netstat`, see Appendix A.

2.3.2 Displaying Socket Statistics

The `socket` subsystem has three attributes that monitor socket listen queue events:

- The `sobacklog_hiwat` attribute counts the maximum number of pending requests to any server socket.
- The `sobacklog_drops` attribute counts the number of times the system dropped a received SYN packet because the number of queued SYN_RCVD connections for a socket equaled the socket's backlog limit.
- The `somaxconn_drops` attribute counts the number of times the system dropped a received SYN packet because the number of queued SYN_RCVD connections for the socket equaled the upper limit on the backlog length (`somaxconn` attribute).

The initial value of these attributes is 0. Use the `sysconfig -q socket` command to display the current attribute values. If the values show that the queues are overflowing, you may need to increase the socket listen queue limit.

The value of the `sominconn` attribute should equal the value of the `somaxconn` attribute. When these two attributes are equal, the value of `somaxconn_drops` will have the same value as `sobacklog_drops`.

However, if the value of the `sominconn` attribute is 0 (the default), and if one or more server applications uses an inadequate value for the backlog argument to its listen system call, the value of `sobacklog_drops` may increase at a rate that is faster than the rate at which the `somaxconn_drops` counter increases. If this occurs, you may want to increase the value of the `sominconn` attribute. See Section 2.2.1.2 for more information.

2.4 Tuning Server Applications

In addition to tuning TCP/IP Services kernel attributes, performance improvements can be made to server applications by:

- Ensuring adequate memory configuration
- Logging IP addresses

2.4.1 Configuring Memory for High Performance

Each connection to an server requires enough memory resources for the following:

- Kernel socket structure
- Internet protocol control block (`inpcb`) structure
- TCP control block structure
- Any socket buffer space that is needed as packets arrive and are consumed

These memory resources total 1 KB for each connection endpoint (not including the socket buffer space), which means 10 MB of memory is required in order to accommodate 10,000 connections.

Your server must have enough memory to handle demanding peak loads. As a rule of thumb, if you configure ten times more memory than the server requires on a busy day, you will have sufficient memory to handle occasional spikes of activity.

Tuning Techniques

2.4 Tuning Server Applications

There are no limitations on a server's ability to handle millions of TCP connections if memory resources are available to service the connections. If memory is insufficient, the server will reject new connection requests until memory is available. Use the `netstat -m` command to monitor the memory that is currently being used by the network subsystem. See Section 2.3.1 for information about displaying memory statistics.

2.4.2 Logging IP Addresses

If your server application logs client host names, the application software may force the system to perform a reverse DNS lookup to obtain the client's host name. Reverse DNS lookups are time-intensive and can cause performance problems on servers with many clients.

Many applications can be modified to log client IP addresses instead of client host names. Logging IP addresses instead of host names may significantly improve the efficiency of the server. Consult the documentation provided by the server software vendor to determine how to disable the logging of client host names.

For example, you can obtain information about modifying Apache HTTP Server software from the Apache HTTP Server documentation site.

2.5 Increasing the Auxiliary Server Connection Limit

The auxiliary server handles a limited number of service invocations in a one-minute period of time. The default is a maximum of 500 connection requests. If the number of requests exceeds this limit, the auxiliary server will not accept additional requests for that service.

If your server receives more than eight requests per second for a service that is spawned by the auxiliary server (for example, POP-3, FTP, and SMTP servers), increase the default connection request limit. You can check the service's log file to determine if a service has been shut down. For example, the file `SYSSYSDEVICE:[TCPIP$POP]TCPIP$POP_RUN.LOG` will contain information about the POP service.

Because the auxiliary server does not spawn any known HTTP server, the connection request limit does not affect HTTP service.

2.6 Solving Performance Problems

This section contains information that you can use to identify and solve server performance problems.

The following tasks can help you to solve performance problems:

- Monitor the server.
See Section 2.3 for information about monitoring the network, the virtual memory subsystem, and network socket statistics.
- Apply any patches recommended for your operating system.
See the Recommended Patch Table for information about operating system patches that can improve performance.
- Apply the kernel attribute values that are recommended for your type of system.
See Section 2.6.1 for a list of attributes that can be tuned to improve performance.

- Prevent “forbidden” and “*url* could not load” messages on Netscape enterprise server systems.

2.6.1 Tuning Recommendations for a Primary Server

This section provides recommendations for tuning a server for optimal performance. These recommendations are applicable to most configurations. The recommendations include the attribute value and a reference to additional information.

The primary recommendations for servers, (including web servers, proxy servers, gateway systems, and firewall systems) are as follows:

- Tune the following socket subsystem attributes:

```
somaxconn = 65535  
sominconn = 65535
```

- Tune the following inet subsystem attributes:

```
pmtu_enabled = 0
```

For proxy servers, gateway systems, and firewall systems, also apply these additional recommendations:

- Modify the following socket subsystem attribute as follows:

```
sbcompress_threshold = 600
```

- Modify the following inet subsystem attribute as follows:

```
ipport_userreserved = 65000
```


A

Troubleshooting Tools Reference

This appendix provides more information about the troubleshooting tools described in this manual. It also describes the tools used for isolating and resolving problems with your network and network software.

To invoke a tool as a command at the system prompt, execute the `SYS$STARTUP:TCPIP$DEFINE_COMMANDS.COM` file. Execution of this file defines each tool as a foreign command.

arp

arp

Displays and controls Address Resolution Protocol (ARP) tables.

Format

```
arp [-u] [-n] hostname
```

```
arp -a [-u] [-n] [-i] hostname
```

```
arp -d hostname
```

```
arp -g hostname
```

```
arp -s [-u] hostname hardware_addr [temp] [pub] [trail]
```

```
arp -f filename
```

Description

The `arp` command displays or modifies the current ARP entry for the host specified by *hostname*. The *hostname* value can be specified by name or IP address, using dotted-decimal notation.

With no flags, the program displays the current ARP entry for *hostname*.

The ARP tables can be displayed by any user, but only privileged users can modify them.

Flags

-a
Displays all current ARP entries.

-d *hostname*
Deletes the entry for *hostname* if the user entering the command is a privileged user.

-f *filename*
Reads entries from *filename* and adds those entries to the ARP tables. Use of this flag requires system privileges. Entries in the file have the following format:

```
hostname hardware_addr [temp] [pub] [trail]
```

Fields in this format are:

Option	Description
<i>hostname</i>	Specifies the remote host identified by the entry.
<i>hardware_addr</i>	Specifies the hardware address of the remote host. The address is given as 6 hexadecimal bytes separated by colons.
temp	Specifies that this ARP table entry is temporary. When this argument is not used, the table entry is permanent.
pub	Indicates that the table entry will be published and that the current system will act as an ARP server, responding to requests for <i>hostname</i> even though the host address is not its own.
trail	Indicates that the trailer encapsulation can be sent to this host.

-g *hostname*

Sends a gratuitous ARP packet. The value for *hostname* can be a local host name, alias, or IP address.

-i *hostname*

Displays the interface with which the ARP entry is associated.

-n *hostname*

Displays numeric IP addresses and hardware addresses only. When this flag is not specified, arp displays hostnames, numeric IP addresses, and hardware addresses.

-s *hostname hardware_addr [temp] [pub]*

Creates a single ARP entry for *hostname*. Use of this flag requires privileges. Fields in the format are:

<i>hostname</i>	Specifies the remote host identified by the entry.
<i>hardware_addr</i>	Specifies the hardware address of the remote given as 6 hexadecimal bytes separated by colons.
temp	Specifies that this ARP table entry is temporary. When this argument is not used, the table entry is permanent.
pub	Indicates that the table entry will be published and the current system will act as an ARP server, responding to requests for <i>hostname</i> even though the host address is not its own.

-u

Displays the MAC address in noncanonical form, with address bytes reversed and separated by a colon character (:). By default, all addresses are displayed in canonical form with address bytes separated by the hyphen character (-).

When used with the *-s* flag, this indicates the *hardware_addr* is specified in noncanonical form.

arp

Examples

The following examples show how to use the arp command.

1. TCPIP> arp -a
a71kt.lkg.dec.com (10.10.2.1) at aa-00-04-00-71-f8 stale
v71kt.lkg.dec.com (10.10.2.3) at aa-00-04-00-70-f8 stale
v72kt.lkg.dec.com (10.10.2.4) at aa-00-04-00-6d-f8
tlab9.lkg.dec.com (10.10.2.11) at aa-00-04-00-42-11
timber.lkg.dec.com (10.10.2.14) at aa-00-04-00-c9-f8

This example shows how to display the ARP address-mapping tables for the local host.

2. TCPIP> arp -a -i
a71kt.lkg.dec.com (10.10.2.1) at aa-00-04-00-71-f8 stale (WE0)
v71kt.lkg.dec.com (10.10.2.3) at aa-00-04-00-70-f8 (WE0)
v72kt.lkg.dec.com (10.10.2.4) at aa-00-04-00-6d-f8 stale (WE0)
tlab9.lkg.dec.com (10.10.2.11) at aa-00-04-00-42-11 (WE0)
timber.lkg.dec.com (10.10.2.14) at aa-00-04-00-c9-f8 (WE0)

This example shows how to display the ARP address-mapping tables for the local host and the interface.

3. TCPIP> arp -s laszlo 08:00:2b:0f:44:23 temp

This example shows how to add a single entry for the remote host laszlo to the ARP mapping tables temporarily. The address is considered canonical even though the bytes are separated by colons. For input, the arp command does not use the colon (:) and hyphen (-) characters to indicate whether the address is canonical or noncanonical. You must have system privileges to execute this command.

4. TCPIP> arp -u -s laszlo 10:00:d4:f0:22:c4 temp

This example shows how to add a single entry for the remote host laszlo to the arp mapping tables temporarily. The -u flag indicates the address is noncanonical. You must have system privileges to execute this command.

5. TCPIP> arp -f newentries

This example shows how to add multiple entries to the ARP mapping tables from a file named newentries. You must have system privileges to execute this command.

dig

Sends domain name query packets to name servers.

Format

```
dig [@server] domain [query-type] [query-class] [+query-option] [-dig-option] [%comment]
```

Description

Domain Information Groper (dig) is a flexible command line tool you can use to gather information from Domain Name System servers. The dig tool has two modes: simple interactive mode, which makes a single query; and batch mode, which executes a query for each query in a list of several query lines. All query options are accessible from the command line.

Parameters

server

Either a domain name or an IP address expressed in dotted-decimal notation. If this optional field is omitted, dig attempts to use the default name server for your machine.

If you specify a domain name, dig resolves the query using the domain name resolver (BIND). If your system does not support DNS, you may have to specify an network address in dotted-decimal notation. Alternatively, if a DNS server is available, that server must be listed in the local hosts database.

domain

The domain name for which you are requesting information. See the `-x` option for a convenient way to specify a reverse translation address query.

query-type

The type of information (DNS query type) that you are requesting. If you omit this parameter, the default value for *query-type* is a (network address). BIND recognizes the following query types:

Query Type	Query Class	Description
a	T_A	Network address
any	T_ANY	All information about the specified domain
mx	T_MX	Mail exchanger for the domain
ns	T_NS	Name servers
soa	T_SOA	Zone of authority record
hinfo	T_HINFO	Host information
axfr	T_AXFR	Zone transfer (must ask an authoritative server)
txt	T_TXT	Arbitrary number of strings (see RFC 1035 for the complete list)

query-class

The network class requested in the query. If you omit this parameter, the default is in (C_IN, Internet class domain). BIND recognizes the following classes:

dig

Query Type	Query Class	Description
in	C_IN	Internet class domain
any	C_ANY	All class information

See RFC 1035 for a complete list of query classes.

You can use the `query-class any` statement to specify a class or a type of query. `dig` parses the first occurrence of `any` to mean `query-type = T_ANY`. To specify `query-class = C_ANY`, you must either specify `any` twice or set `query-class` using the `-c` option.

Options

%ignored-comment

Use the percent (%) character to include an argument that is not parsed. This can be useful if you are running `dig` in batch mode. Instead of resolving every `@server-domain-name` in a list of queries, you can avoid the overhead of doing so, and still have the domain name on the command line as a reference. For example:

```
dig @128.9.0.32 %venera.isi.edu mx isis.edu
```

-<dig-option>

Use the hyphen (-) character to specify an option that affects the operation of `dig`. The options described in the Table A-1 are currently available (although not guaranteed to be useful). Options that are uppercase characters must be specified in quotes. For example, `dig -"P"`

Table A-1 dig Options

Option	Description
<code>-x dot-notation-address</code>	Convenient form for specifying reverse translation of IP address. Instead of: <pre>dig 32.0.9.128.in-addr.arpa</pre> you can use: <pre>dig -x 128.9.0.32</pre>
<code>-f file</code>	File for <code>dig</code> batch mode. The file contains a list of query specifications (<code>dig</code> command lines) that are to be executed successively. Lines beginning with <code>;</code> , <code>#</code> , or <code>\n</code> are ignored. Other options can still appear on the command line and will be in effect for each batch query.
<code>-T time</code>	Time (in seconds) between start of successive queries when running in batch mode. Can be used to keep two or more batch <code>dig</code> commands running synchronously. The default is 0.
<code>-p port</code>	Port number. Queries a name server listening to a nonstandard port number. The default is 53.

(continued on next page)

Table A-1 (Cont.) dig Options

Option	Description
-P	After query returns, executes a ping command to compare response times. This option issues the following command: \$ MCR TCPIP\$PING -C 3 <i>server_name</i>
-t <i>query-type</i>	Type of query. Specifies either an integer value to be included in the type field, or uses the abbreviated mnemonic (such as mx).
-c <i>query-class</i>	Class of query. Specifies either an integer value to be included in the class field, or use the abbreviated mnemonic (such as in).

+<query-option>

Use the plus (+) character to specify an option to be changed in the query packet or to change dig output specifics. Many of these options are the same options accepted by nslookup. If an option requires a parameter, use the following format:

+ *keyword* [=value]

Most keywords can be abbreviated. Parsing of the "+" options is very simplistic—a value must not be separated from its keyword by any spaces. The following keywords are currently available:

Keyword	Abbreviation	Default	Description
[no] debug	deb	deb	Turn on/off debugging mode.
[no] d2		nod2	Turn on or off extra debugging mode.
[no] recurse	rec	rec	Use or do not use recursive lookup.
retry=#	ret	4	Set number of retries to #.
time=#	ti	4	Set timeout length to # seconds.
[no] ko		noko	Keep open option (implies vc).
[no] vc		novc	Use or do not use virtual circuit.
no defname	def	def	Use or do not use default domain name.
[no] search	sea	sea	Use or do not use domain search list.
domain=NAME	do		Set default domain name to NAME.
[no]ignore	i	noi	Ignore or do not ignore truncation errors.
[no] primary	pr	nopr	Use or do not use primary server.
no aaonly	aa	noaa	Authoritative query only flag.
[no] cmd		cmd	Echo parsed arguments.
[no] stats	st	st	Display query statistics.

dig

Keyword	Abbreviation	Default	Description
[no] Header	H	H	Display basic header.
[no] header	he	he	Display header flags.
[no] ttlid	tt	tt	Display TTLs.
[no] cl		nocl	Display class information.
[no] qr		noqr	Display outgoing query
[no] reply	rep	rep	Display reply.
[no] ques	qu	qu	Display question section.
[no] answer	an	an	Display answer section.
[no] author	au	au	Display authoritative section.
[no] addit	ad	ad	Display additional section.
pfdef			Set to default display flags.
pfmin			Set to minimal default display flags.
pfset=#			Set display flags to # (# can be hexadecimal, octal, or decimal).
pfand=#			Bitwise and display flags with #.
pfor=#			Bitwise or display flags with #.

Examples

The following examples show how to use the dig command.

```
1. $ dig
; <<>> DiG 8.1 <<>>
;; res options: init recurs defnam dnsrch
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 6
;; flags: qr rd ra; QUERY: 1, ANSWER: 13, AUTHORITY: 0, ADDITIONAL: 13
;; QUERY SECTION:
;;      ., type = NS, class = IN
;; ANSWER SECTION:
.          1d20h1m11s IN NS  E.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  D.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  A.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  H.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  C.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  G.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  F.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  B.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  J.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  K.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  L.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  M.ROOT-SERVERS.NET.
.          1d20h1m11s IN NS  I.ROOT-SERVERS.NET.
```

```

;; ADDITIONAL SECTION:
E.ROOT-SERVERS.NET.      2d20h1m11s IN A  192.203.230.10
D.ROOT-SERVERS.NET.      2d20h1m11s IN A  128.8.10.90
A.ROOT-SERVERS.NET.      2d20h1m11s IN A  198.41.0.4
H.ROOT-SERVERS.NET.      2d20h1m11s IN A  128.63.2.53
C.ROOT-SERVERS.NET.      2d20h1m11s IN A  192.33.4.12
G.ROOT-SERVERS.NET.      2d20h1m11s IN A  192.112.36.4
F.ROOT-SERVERS.NET.      2d20h1m11s IN A  192.5.5.241
B.ROOT-SERVERS.NET.      2d20h1m11s IN A  128.9.0.107
J.ROOT-SERVERS.NET.      2d20h1m11s IN A  198.41.0.10
K.ROOT-SERVERS.NET.      2d20h1m11s IN A  193.0.14.129
L.ROOT-SERVERS.NET.      2d20h1m11s IN A  198.32.64.12
M.ROOT-SERVERS.NET.      2d20h1m11s IN A  202.12.27.33
I.ROOT-SERVERS.NET.      2d20h1m11s IN A  192.36.148.17

;; Total query time: 4013 msec
;; FROM: lassie.ucx.lkg.dec.com to SERVER: default -- 16.20.208.53
;; WHEN: Wed Aug  9 16:42:08 2000
;; MSG SIZE sent: 17 rcvd: 436

```

This example shows how to query your default name server for query type NS (default query type) and query class IN (default query class). The output shows the address records for the root name servers and their IP addresses.

2. \$ dig microsoft.com mx

```

;<<>> DiG 8.1 <<>> microsoft.com mx
;; res options: init recurs defnam dnsrch
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 6
;; flags: qr rd ra; QUERY: 1, ANSWER: 5, AUTHORITY: 12, ADDITIONAL: 8
;; QUERY SECTION:
;;      microsoft.com, type = MX, class = IN

;; ANSWER SECTION:
microsoft.com.      2h18m8s IN MX  10 mail2.microsoft.com.
microsoft.com.      2h18m8s IN MX  10 mail3.microsoft.com.
microsoft.com.      2h18m8s IN MX  10 mail4.microsoft.com.
microsoft.com.      2h18m8s IN MX  10 mail5.microsoft.com.
microsoft.com.      2h18m8s IN MX  10 mail1.microsoft.com.

;; AUTHORITY SECTION:
com.      5d22h12m9s IN NS  A.ROOT-SERVERS.NET.
com.      5d22h12m9s IN NS  E.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  F.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  F.ROOT-SERVERS.NET.
com.      5d22h12m9s IN NS  J.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  K.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  A.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  M.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  G.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  C.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  I.GTLD-SERVERS.NET.
com.      5d22h12m9s IN NS  B.GTLD-SERVERS.NET.

;; ADDITIONAL SECTION:
mail2.microsoft.com. 2h8m41s IN A  131.107.3.124
mail3.microsoft.com. 2h27s  IN A  131.107.3.123
mail4.microsoft.com. 1h53m4s IN A  131.107.3.122
mail5.microsoft.com. 2h8m43s IN A  131.107.3.121
mail1.microsoft.com. 2h8m43s IN A  131.107.3.125
A.ROOT-SERVERS.NET. 2d19h47m37s IN A  198.41.0.4
E.GTLD-SERVERS.NET. 1d9h45m57s IN A  207.200.81.69
F.GTLD-SERVERS.NET. 3h16m16s  IN A  198.17.208.67

```

dig

```
;; Total query time: 4019 msec  
;; FROM: lassie.ucx.lkg.dec.com to SERVER: default -- 16.20.208.53  
;; WHEN: Wed Aug 9 16:55:42 2000  
;; MSG SIZE sent: 31 rcvd: 493
```

This example shows how to obtain the mail server records for Microsoft.

ifconfig

Assigns an address to a network interface, and configures and displays network interface parameters.

Format

```
ifconfig interface_id [address_family] [address[/bitmask]] [dest_address] [parameters])
```

```
ifconfig -a [-d] [-u] [-v] [address_family]
```

```
ifconfig -l [-d] [-u] [-v] [address_family]
```

```
ifconfig [-v] interface-id [address_family]
```

Description

Use the `ifconfig` command to define the network address of each interface. You can also use the `ifconfig` command at other times to display all interfaces that are configured on a system, to redefine the address of an interface, or to set other operating parameters.

Note

If you want to redefine the interface address or the net mask, you should stop TCP/IP Services. Otherwise, any TCP/IP process currently running will continue to use the old address and net mask and will fail.

Any user can query the status of a network interface; only a privileged user can modify the configuration of network interfaces.

You specify an interface with the `ifconfig interface-id` syntax. (See your hardware documentation for information on obtaining an interface ID.)

If you specify only *interface-id*, the `ifconfig` program displays the current configuration for the specified network interface only.

If a protocol family is specified by the *address_family* parameter, `ifconfig` reports only the configuration details specific to that protocol family.

When changing an interface configuration, if the address family cannot be inferred from the *address* parameter, an address family must be specified. The address family is required because an interface can receive transmissions in different protocols, each of which can require a separate naming scheme.

The *address* parameter is the network address of the interface being configured. For the `inet` address family, the *address* parameter is either a host name or an IP address in the standard dotted-decimal notation with or without the optional Classless Inter-Domain Routing (CIDR) bit mask (*bitmask*). If you specify *bitmask*, do not use the *netmask* parameter.

The destination address (*dest_address*) parameter specifies the address of the correspondent on the remote end of a point-to-point link.

ifconfig

Flags

- a**
Displays information about all interfaces that are configured on a system.
- d**
Displays information about interfaces that are down.
- l**
Displays interface names that are configured on a system.
- u**
Displays information about interfaces that are up.
- v**
Displays detailed information about interfaces, such as hardware addresses.

Parameters

alias *alias_address*[/*bitmask*]

Establishes an additional network address for this interface. This can be useful when changing network numbers and you want to continue to accept packets addressed to the old interface.

If you do not specify a bit mask or net mask with the alias address, the default net mask is based on the alias address's network class.

If you are using the optional bit mask argument, do not use the net mask argument.

-alias *alias_address*

Removes the network address specified. This can be used if you incorrectly specified an alias or if an alias is no longer needed. The `-alias` parameter functions in the same manner as the delete parameter.

aliaslist *address_list*[/*bitmask*]

Establishes a range of additional network addresses for this interface. The range can be a comma-separated list or a hyphenated list, and is inclusive. You can also specify the optional CIDR bit mask (*bitmask*) argument at the end of the list. Do not use both a comma-separated list and a hyphenated list for a range.

-aliaslist

Removes a range of network addresses for this interface. This can be useful when deleting network numbers and you want to keep the primary interface address. The `-aliaslist` rules are the same as for the `aliaslist` parameter.

allmulti

Enables the reception of all multicast packets.

-allmulti

Disables the reception of all multicast packets.

arp

Enables the use of the Address Resolution Protocol (ARP) in mapping between network-level addresses and link-level addresses. This parameter is on by default.

-arp

Disables the use of the ARP. Use of this parameter is not recommended.

broadcast *broad_address*

Specifies the address to use to represent broadcasts to the network. The default broadcast address is the address with a host part consisting of all ones (1).

Note that the computation of the host part is dependent on *netmask* (see the description of the *netmask* parameter).

delete [*net_address*]

Removes the network address specified. Use this parameter if you incorrectly specified an alias or if the alias was no longer needed. If you have incorrectly set an NS address, removing all NS addresses will allow you to specify the host portion again.

If no address is specified, *ifconfig* deletes all network addresses for the interface.

down

Marks an interface as not working (down), which keeps the system from trying to transmit messages through that interface. If possible, the *ifconfig* command also resets the interface to disable the reception of messages. Routes that use the interface, however, are not automatically disabled.

ipmtu *mtu_value*

Alters the size of the maximum transmission unit (MTU) for messages that your system transmits. It might be necessary to reduce the MTU size so that bridges connecting token rings can transfer frames without error.

metric *number*

Sets the routing metric, or number of hops, for the interface to the value of *number*. The default value is zero (0) if *number* is not specified, indicating that both hosts are on the same network. The routing metric is used by the ROUTED and GATED services, with higher metrics indicating that the route is less favorable.

netmask *mask*

Specifies how much of the address to reserve for subdividing networks into sub-networks. This parameter can only be used with an address family of *inet*. Do not use this parameter if you are specifying the CIDR mask (*/bitmask*) with the *address* parameter, *alias* parameter, or *aliaslist* parameter.

The *mask* variable includes both the network part of the local address and the subnet part, which is taken from the host field of the address. The mask can be specified as a single hexadecimal number beginning with 0x, in the standard Internet dotted-decimal notation, or beginning with a name.

The mask contains ones (1) for the bit positions in the 32-bit address that are reserved for the network and subnet parts, and zeros (0) for the bit positions that specify the host. The mask should contain at least the standard network portion.

The default net mask is based on the address parameter's network class.

up

Marks an interface as working (up). This parameter is used automatically when setting the first address for an interface, or it can be used to enable an interface after an *ifconfig down* command. If the interface was reset when previously marked with the *down* parameter, the hardware will be reinitialized.

ifconfig

Examples

The following examples show how to use the `ifconfig` command.

1. TCPIP> `ifconfig sl0`
`sl0: flags=10`

This example shows how to query the status of serial line interface `sl0`.

2. TCPIP> `ifconfig lo0 inet 127.0.0.1 up`

This example shows how to configure the local loopback interface. Only a user with system privileges can modify the configuration of a network interface.

3. TCPIP> `ifconfig ln0 212.232.32.1/22`

This example shows how to configure an `ln0` interface. The broadcast address is `212.232.35.255` as the 22-bit mask specifies four Class C networks.

4. TCPIP> `ifconfig tra0 130.180.4.1/24 speed 4`

This example shows how to configure the token ring interface for a 4 Mb/s token ring with a net mask of `255.255.255.0` in CIDR format.

5. TCPIP> `ifconfig tra0 down`
TCPIP> `ifconfig tra0 speed 16 up`

This example shows how to stop the token ring interface and start it for a 16 Mb/s token ring.

6. TCPIP> `ifconfig we0 alias 132.50.40.35/24`

This example shows how to add alias `132.50.40.35` with a net mask of `255.255.255.0` in CIDR format to interface `we0`.

7. TCPIP> `ifconfig we0 aliaslist 132.240.32-36.40-50/24`

This example shows how to add network addresses 40 through 50, to subnets `18.240.32`, `18.240.33`, `18.240.34`, `18.240.35`, and `18.240.36` with a net mask of `255.255.255.0` in CIDR format to the `we0` interface.

8. TCPIP> `ifconfig we0 down delete abort`
`145.92.16.1: aborting 7 tcp connection(s)`

This example shows how to stop Ethernet interface `we0`, delete all addresses associated with the interface, and close all TCP connections.

9. TCPIP> `ifconfig we0 -alias 145.92.16.2 abort`
`145.92.16.2: aborting 2 tcp connection(s)`

This example shows how to delete the alias address `145.92.16.2` on interface `tu0` and close all TCP connections.

10. TCPIP> `ifconfig we0 alias 145.92.16.2 physaddr aa:01:81:43:02:11`

This example shows how to associate MAC address `aa:01:81:43:02:11` with the alias address `145.92.16.2`.

11. TCP/IP> ifconfig we0 -alias 145.92.16.2 -physaddr aa:01:81:43:02:11

This example shows how to disassociate MAC address aa:01:81:43:02:11 from the alias address 145.92.16.2.

12. A72KT: ifconfig -l
TCP/IP> ifconfig -l
LO0 TN0 WE0

This example shows how to display the names of the interfaces on the system only.

13. TCP/IP> ifconfig -v WE0
we0: flags=c43<UP,BROADCAST,RUNNING,MULTICAST,SIMPLEX>
HWaddr aa:0:4:0:72:f8
inet 10.10.2.2 netmask ffffffff broadcast 10.10.2.255 ipmtu 1500

This example shows how to display the hardware and IP addresses of interface WE0.

ndc

ndc

Manages the BIND server.

Format

```
ndc directive [directive ... ]
```

Description

This command allows the name server administrator to send various messages to a name server. You can specify zero or more directives from the following list.

Directives

status

Displays the current status of the BIND server process.

dumpdb

Causes the BIND server to dump its database and cache to SYSSSPECIFIC:[TCPIP\$BIND]TCPIP\$BIND_SERVER_ZONES_DUMP.DB.

reload

Causes the BIND server to check the serial numbers of all primary and secondary zones and to reload those that have changed.

stats

Causes the BIND server to dump statistics to SYSSSPECIFIC:[TCPIP\$BIND]TCPIP\$BIND_SERVER_STATISTICS.LOG.

trace

Causes the BIND server to increment its tracing level by 0. Trace information is written to SYSSSPECIFIC:[TCPIP\$BIND]TCPIP\$BIND_RUN.LOG. Higher tracing levels result in more detailed information.

notrace

Causes the BIND server to set its tracing level to 0.

start

Causes the BIND server to be started, if it is not running.

stop

Causes the BIND server to be stopped if it is running.

restart

Causes the BIND server to be stopped and restarted.

Examples

The following examples show how to use the ndc command.

1. `$ ndc status`
BIND Server process information:
Process ID: 44C0021C
Process name: TCPIP\$BIND_1
Priority: 9
Elapsed CPU time: 0 00:00:31.19
Buffered I/O count: 214082
Direct I/O count: 404
Page Faults: 485
Pages: 4096
Peak virtual size: 173696
Peak working set size: 5920
Process state: LEF

This example shows how to display the current status of the BIND server process.

2. `$ ndc dumpdb`

This example shows how to dump the BIND server's database into the `SYSSSPECIFIC:[TCPIP$BIND]TCPIP$BIND_SERVER_ZONES_DUMP.DB` file. Use the DCL command TYPE to view the contents of this file.

3. `$ ndc stats`

This example shows how to dump BIND server statistics to the `SYSSSPECIFIC:[TCPIP$BIND]TCPIP$BIND_SERVER_ZONES_STATISTICS.LOG` file. Use the DCL command TYPE to view the contents of this file.

netstat

netstat

Displays network-related data in various formats.

Format

```
netstat [-rn | [-an] [-f address_family] [interval]
```

```
netstat [-adHimMnrstv] [-f address_family] [interval]
```

```
netstat [-ntdz] [-I interface] [interval]
```

Description

The interval argument specifies in seconds the interval for updating and displaying information. The first line of the display shows cumulative statistics; subsequent lines show statistics recorded during interval.

Default display

When used without flags, the `netstat` command displays a list of active sockets for each protocol. The default display shows the following items:

- Local and remote addresses
- Send and receive queue sizes (in bytes)
- Protocol
- State

Address formats are of the form *host.port* or *network.port* if a socket's address specifies a network but no specific host address. The host and network address are displayed symbolically unless `-n` is specified.

Interface display

The network interface display provides a table of cumulative statistics for the following:

- Interface name (Name)
- Maximum transmission unit (Mtu)
- Network address
- Packets received (Ipkts)
- Packets received in error (Ierrs)
- Packets transferred (Opkts)
- Outgoing packets in error (Oerrs)
- Collisions (Coll)

Note that the collisions item has different meanings for different network interfaces.

- Drops (optional with `-d`)
- Timers (optional with `-t`)

Routing table display

A route consists of a destination host or network and a gateway to use when forwarding packets. Direct routes are created automatically for each interface attached to the local host when you issue the `ifconfig` command. Routes can be modified automatically in response to the prevailing condition of the network.

The routing-table display format indicates available routes and the status of each in the following fields:

Flags

Displays the state of the route as one or more of the following:

U	Up, or available.
G	This route is to a gateway.
H	This route is to a host.
D	This route was dynamically created by a redirect.
M	This route was modified by a redirect.
S	This is a static route that was created by the <code>route</code> command.
R	This is a reject route that was created by the <code>route</code> command.

refcnt

Gives the current number of active uses for the route. Connection-oriented protocols hold on to a single route for the duration of a connection; connectionless protocols obtain routes in the process of sending to a destination.

use

Provides a count of the number of packets sent using the route.

interface

Indicates the network interface used for the route.

When the `-v` flag is specified, the routing table display includes the route metrics. An asterisk (*) indicates the metric is locked.

Flags**-a**

Displays the state of sockets related to the Internet protocol. Includes sockets for processes such as servers that are currently listening at a socket but are otherwise inactive.

-d

Displays the number of dropped packets; for use with the `-I` interface or `-i` flags. You can also specify an interval argument (in seconds).

-f *address_family*

Limits reports to the specified address family. The address families that can be specified might include the following:

inet	Specifies reports of the AF_INET family, if present in the kernel.
all	Lists information about all address families in the system.
any	Lists information about any address families in the system.

-H

Displays the current ARP table (behaves like `arp -a`).

netstat

-i

Displays the state of configured interfaces. (Interfaces that are statically configured into the system but not located at system startup are not shown.)

When used with the `-a` flag, `-i` displays IP and link-level addresses associated with the interfaces.

You can use the `-i` flag to retrieve your system's hardware address.

-I interface

Displays information about the specified interface.

-m

Displays information about memory allocated to data structures associated with network operations.

-M

Displays Internet protocol multicast routing information. When used with the `-s` flag, `-M` displays IP multicast statistics.

-n

Displays network address in numerical format with network masks in CIDR format. When this flag is not specified, the address is displayed as host name and port number. This flag can be used with any of the display formats.

-r

Displays the host's routing tables. When used with the `-s` flag, `-r` shows the host's routing statistics instead of its routing tables.

-s

Displays statistics for each protocol.

-t

Displays timer information. Use with the `-I` interface or `-i` flag.

-v

Displays detailed output when specified with the `-r` flag. In this case, route metric values are displayed.

-z

Sets the network interface counters to zero. This flag must be specified with the `-I` interface flag. You must have system privileges to use the `-z` flag.

Examples

The following examples show how to use the `netstat` command to display information about configured interfaces and routing tables.

1. TCPIP> netstat -i

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
LO0	4096	<Link>	Link#1	167	0	167	0	0
LO0	4096	loop	LOCALHOST	167	0	167	0	0
SE0	1500	<Link>	aa:0:4:0:6d:f8	1544	0	157	0	0
SE0	1500	loop	RUFUS	1544	0	157	0	0
TN0*	1280	<Link>	Link#2	0	0	0	0	0

This example shows how to display the state of the configured interfaces.

```

2. TCPIP> netstat -r
Routing tables
Destination      Gateway          Flags    Refs    Use  Interface
Route Tree for Protocol Family 26:
Route Tree for Protocol Family 2:
default          sqagate          UG       0        0  SE0
10.10.2          v72kt            U        2       125 SE0
v72kt           v72kt            UHL      0        50 SE0
LOCALHOST       LOCALHOST        UHL      7       117 LO0

```

This example shows how to display the routing tables.

```

3. TCPIP> netstat -rn
Routing tables
Destination      Gateway          Flags    Refs    Use  Interface
Route Tree for Protocol Family 26:
Route Tree for Protocol Family 2:
default          10.10.2.66       UG       0        0  SE0
10.10.2/24       10.10.2.4        U        2       109 SE0
10.10.2.4        10.10.2.4        UHL      0        50 SE0
127.0.0.1        127.0.0.1        UHL      7       117 LO0

```

This example shows how to display the routing tables with network addresses.

nslookup

nslookup

Queries Internet name servers interactively.

Format

```
nslookup [-option ...] [host_to_find | - [server] ]
```

Description

The `nslookup` command is a program that is used to query Internet domain name servers. The `nslookup` command has two modes: noninteractive and interactive.

Noninteractive mode

Noninteractive mode is used to display just the name and requested information for a host or domain. Noninteractive mode is invoked when the name or Internet address of the host to be looked up is given as the first argument. The optional second argument specifies the host name or address of a name server.

Interactive mode

Interactive mode allows the user to query name servers for information about various hosts and domains or to display a list of hosts in a domain. Interactive mode is invoked when you specify `nslookup` without arguments (the default name server will be used), or when the first argument you specify is a hyphen (-) and the second argument is the host name or IP address of a name server.

The options listed under the `set` command can be specified in the `.nslookuprc` file in the user's home directory if they are listed one per line. Options can also be specified on the command line if they precede the arguments and are prefixed with a hyphen (-). For example, to change the default query type to host information, and the initial time out to 10 seconds, enter the following command:

```
$ nslookup -query=hinfo -timeout=10
```

Interactive commands

Commands can be interrupted at any time by pressing `Ctrl/C`. To exit, press `Ctrl/D` (EOF) or type `exit`. The command line length must be less than 256 characters. To treat a built-in command as a host name, prefix it with an escape character (^) plus a backslash (\). Note that an unrecognized command will be interpreted as a host name.

Commands

***host* [server]**

Looks up information for the host using either the current default server or the specified server. If `host` is an IP address and the query type is `A` or `PTR`, the name of the host is returned. If `host` is a name and does not have a trailing period, the default domain name is appended to the name. (This behavior depends on the state of the `set` options `domain`, `srchlist`, `defname`, and `search`.) To look up a host not in the current domain, append a dot (.) to the end of the domain name.

server *domain***lserver *domain***

Changes the default server to *domain*. The `lserver` command uses the initial server to look up information about *domain*, while the `server` command uses the current default server. If an authoritative answer cannot be found, the names of servers that might have the answer are returned.

root

Changes the default server to the server for the root of the domain name space. Currently, the host `ns.internic.net` is used. (This command is a synonym for `lserver ns.internic.net`.) The name of the root server can be changed with the `set root` command.

finger [*name*] [> *filename*]**finger [*name*] [>> *filename*]**

Connects with the `finger` server on the current host. The current host is defined when a previous lookup for a host was successful and returned address information (see the `set querytype=A` command). The redirection symbols (`>` and `>>`) can be used to redirect output in the usual manner.

ls [*option*] *domain* [> *filename*]**ls [*option*] *domain* [>> *filename*]**

Lists the information available for *domain*, optionally creating or appending to *filename*. The default output contains host names and their IP addresses. The value for *option* can be one of the following:

Option	Description
<code>-t <i>querytype</i></code>	Lists all records of the specified type. (See <i>querytype</i> in Table A-2.)
<code>-a</code>	Lists aliases of hosts in the domain. This option is a synonym for <code>-t CNAME</code> .
<code>-d</code>	Lists all records for the domain. This option is a synonym for <code>-t ANY</code> .
<code>-h</code>	Lists CPU and operating system information for the domain. This option is a synonym for <code>-t HINFO</code> .
<code>-s</code>	Lists well-known services of hosts in the domain. This option is a synonym for <code>-t WKS</code> . When output is directed to a file, a pound sign (#) is displayed for every 50 records received from the server.

view *filename*

Sorts and lists the output of previous `ls` commands.

help

Displays a brief summary of commands.

exit

Exits the program.

set *keyword* [=*value*]

Use the command to change state information that affects the lookups. Table A-2 lists the valid keywords.

Table A-2 Options to the nslookup set Command

Keyword	Function
ALL	<p>Displays the current values of the options you can set as well as information about the current default server. For example:</p> <pre>> set all</pre>
class=value	<p>Changes the query class to one of the following:</p> <ul style="list-style-type: none"> • IN — The internet class (default) • CHAOS — The chaos class • ANY — Wildcard <p>The class specifies the protocol group of the information. You abbreviate this keyword to <code>cl</code>.</p> <p>This command tells nslookup to resolve both <code>in</code> and <code>chaos</code> class queries (you can enter <code>in</code> and <code>chaos</code>):</p> <pre>> set class=ANY</pre>
querytype	<p>Specifies the type of information you want. For example:</p> <pre>> set querytype=A > set querytype=ANY</pre> <p>Valid types are:</p> <p>SOA Start of authority. Marks the beginning of a zone's data and defines parameters that affect the entire zone.</p> <p>NS Name server. Identifies a domain's name server.</p> <p>A Address. Maps a host name to an address.</p> <p>ANY Defines all available resource records for a given name.</p> <p>PTR Pointer. Maps an address to a host name.</p> <p>MX Identifies where to deliver mail for a given domain.</p> <p>CNAME Defines an alias host name.</p> <p>HINFO Host information. Describes a host's hardware and operating system.</p> <p>WKS Well-known service. Advertises network services.</p>
[no]debug	<p>Turns on debugging (default is <code>nodebug</code>). nslookup displays detailed information about the packet sent to the server and the answer. For example:</p> <pre>> set debug</pre> <p>You can use the abbreviations <code>nodeb</code> and <code>deb</code>.</p>

(continued on next page)

Table A–2 (Cont.) Options to the nslookup set Command

Keyword	Function
[no]d2	Returns all-inclusive debugging information (default is nod2). Displays all the fields of every packet. For example: > set d2
recurse	Tells the BIND server to contact other servers if it does not have the information you want. The servers carry out a complete (recursive) resolution for each query. For example: > set recurse
retry	Number of times that nslookup attempts to contact a BIND server if repeated tries fail. For example: > set retry=8
timeout	Length of time to wait for a reply from each attempt. For example: > set timeout=9
root=value	Changes the root server. For example, the following command changes the root server to ns.nasa.gov. > set root=ns.nasa.gov
ignoretc	Tells nslookup to ignore packet truncation errors. For example: > set ignoretc
domain name	Changes the default domain to the domain you specify. The settings of the defname and search options control how the default domain name is appended to lookup requests. The domain search list contains the parents of the default domain if the default domain has at least two components in its name. The default value is set in the TCP/IP configuration database. To specify the default, type the abbreviation do. For example, if the default domain is CC.Berkeley.EDU, the search list is CC.Berkeley.EDU and Berkeley.EDU.
srchlist	If set, nslookup appends each of the domain names specified in the srchlist option to an unqualified host name and performs a query until an answer is received.

(continued on next page)

nslookup

Table A–2 (Cont.) Options to the nslookup set Command

Keyword	Function
<code>srchlist=names</code>	<p>Changes the default domain name to the first name you specify, and changes the domain search list to all the names you specify. Specify a maximum of six names separated by slashes (/).</p> <p>In the following example, the command sets the default domain to <code>lcs.MIT.EDU</code> and changes the search list to the three specified domains. The command overrides the default domain name and associated search list for the <code>set domain</code> command.</p> <pre>> set srchlist=lcs.MIT.EDU/ai.MIT.EDU/MIT.EDU</pre> <p>The default is the domain name specified in the TCP/IP configuration database. The abbreviated form of the command is <code>srchl</code>.</p>
<code>[no]defname</code>	<p>Tells nslookup to append a default domain name to a lookup request if the specified DNS name is not fully qualified.¹ The abbreviated form is <code>[no]def</code>.</p> <p>For example, an nslookup query for the host <code>rainy</code> becomes <code>rainy.cc.berkeley.edu</code>.</p>
<code>[no]search</code>	<p>Tells nslookup to append the search list domain names to the lookup request domain name if the lookup request domain name is not fully qualified.¹ The default is <code>search</code>. The abbreviated form is <code>[no]sea</code>.</p>

¹A fully qualified domain name is a name that ends with a dot (`.`), as in *host.domain*.

Examples

The following example shows how to use nslookup interactively.

- ```
$ nslookup
Default Server: condor.lgk.dec.com
Address: 16.99.208.53

> set all
Default Server: condor.lgk.dec.com
Address: 16.99.208.53

Set options:
nodebug defname search recurse
nod2 novc noignoretc port=53
querytype=A class=IN timeout=4 retry=4
root=a.root-servers.net.
domain=xyz.prq.dec.com
srchlist=xyz.prq.dec.com
```

---

## ping

Send ICMP ECHO\_REQUEST packets to network hosts.

### Format

```
ping [-dfnqrurvR] -c count [-i wait] [-l preload] [-p pattern] [-s packetsize] host
```

### Description

The `ping` command uses the ICMP (Internet Control Message Protocol) mandatory ECHO\_REQUEST datagram to elicit an ICMP ECHO\_RESPONSE message from the specified host or gateway host. ECHO\_REQUEST datagrams (`ping`) have an IP (Internet Protocol) and ICMP header, followed by a `struct timeval` and then an arbitrary number of pad bytes used to fill out the packet.

When using `ping` for fault isolation, first run the command on the local host to verify that the local network interface is up and running. Then, hosts and gateways further and further away should be sent the `ping` command. Round-trip times and packet loss statistics are computed. If duplicate packets are received, they are not included in the packet loss calculations, although the round-trip time of these packets is used in calculating the minimum, average, and maximum round-trip time numbers. When the specified number of packets have been sent (and received), or if the program is terminated with a SIGINT, a brief summary is displayed.

This program is intended for use in network testing, measurement, and management. Because of the load it can impose on the network, it is unwise to use `ping` during normal operations or from automated scripts.

#### ICMP packet details

An IP header without options is 20 bytes. An ICMP ECHO\_REQUEST packet contains an additional 8 bytes worth of ICMP header followed by an arbitrary amount of data. When a `packetsize` is given, this indicates the size of this extra piece of data (the default is 56). Thus, the amount of data received inside of an IP packet of type ICMP ECHO\_REPLY will always be 8 bytes more than the requested data space (the ICMP header).

If the data space is at least 8 bytes large, `ping` uses the first 8 bytes of this space to include a timestamp, which it uses in the computation of round-trip times. If less than 8 bytes of pad are specified, no round-trip times are given.

#### Duplicate and damaged packets

The `ping` command will report duplicate and damaged packets. Duplicate packets should never occur, and seem to be caused by inappropriate link-level retransmissions. Duplicates can occur in many situations and are rarely (if ever) a good sign, although the presence of low levels of duplicates can not always be cause for alarm.

Damaged packets are obviously serious cause for alarm and often indicate broken hardware somewhere in the `ping` packet's path (in the network or in the hosts).

## ping

### Different data patterns

The network layer should never treat packets differently depending on the data contained in the data portion. Unfortunately, data-dependent problems have been known to invade networks and remain undetected for long periods of time. In many cases the problematic pattern does not have sufficient transitions, such as all ones (1) or all zeros (0), or has a pattern at the right, such as almost all zeros (0). It is not necessarily enough to specify a data pattern of all zeros on the command line because the problematic pattern of interest is at the data-link level, and the relationship between what you enter and what the controllers transmit can be complicated.

Data-dependent problems can be identified only by extensive testing. If you are lucky, you can manage to find a file that either cannot be sent across your network or that takes much longer to transfer than other files of similar length. You can then examine this file for repeated patterns that you can test by using the `-p` option to the `ping` command.

### TTL details

The TTL value of an IP packet represents the maximum number of IP routers that the packet can go through before being thrown away. You can expect each router in the Internet to decrement the TTL field by exactly one.

The maximum possible value of this field is 255, and most UNIX compatible systems set the TTL field of ICMP ECHO\_REQUEST packets to 255. This is why you can use the `ping` command on some hosts but not reach them with TELNET or FTP.

In normal operation, `ping` displays the TTL value from the packet it receives. When a remote system receives a `ping` packet, it can do one of three things with the TTL field in response:

- Not change the field. This is what Berkeley UNIX compatible systems did before BSD Version 4.3. In this case, the TTL value in the received packet will be 255 minus the number of routers in the round-trip path.
- Set the field to 255. This is what current Berkeley UNIX compatible systems do. In this case, the TTL value in the received packet will be 255 minus the number of routers in the path from the remote system to the host that received the `ping` commands.
- Set the field to some other value. Some machines use the same value for ICMP packets that they use for TCP packets; for example, either 30 or 60. Others may use completely wild values.

## Cautions

Many hosts and gateways ignore the `RECORD_ROUTE` option.

Flooding and preloading the `ping` command is generally not recommended, and flooding `ping` messages on the broadcast address should be done only under very controlled conditions.

## Flags

### **-c count**

Stops after sending (and receiving) the specified number (*count*) of ECHO\_RESPONSE packets.

**-d**

Sets the `SO_DEBUG` option on the socket being used.

**-f**

Floods ping. Outputs packets as fast as they come back or 100 times per second, whichever is more. For every `ECHO_REQUEST` sent, a dot (.) is displayed, while for every `ECHO_REPLY` received a backspace is used. This provides a rapid display of how many packets are being dropped. You must have system privileges to use this option. Using the `-f` flag can be very hard on a network and should be used with caution.

**-i wait**

Waits the specified number of seconds between sending each packet. The default is to wait for 1 second between each packet. This option is incompatible with the `-f` option.

**-l preload**

If `preload` is specified, ping sends that many packets as fast as possible before falling into its normal mode of behavior. You must have system privileges to use this option. Using the `-l` option can be very hard on a network and should be used with caution.

**-n**

Numeric output only. No attempt is made to look up symbolic names for host addresses. This occurs only when displaying ICMP packets other than `ECHO_RESPONSE`.

**-p pattern**

Specifies up to 16 pad bytes to fill out the packet you send. This is useful for diagnosing data-dependent problems in a network. For example, `-p ff` will cause the sent packet to be filled with all ones (1).

**-q**

Suppresses output. Nothing is displayed except the summary lines at startup time and at completion.

**-R**

Records route. Includes the `RECORD_ROUTE` option in the `ECHO_REQUEST` packet and displays the route buffer on returned packets. Note that the IP header is large enough for only nine such routes. Many hosts ignore or discard this option.

**-r**

Bypasses the normal routing tables and sends directly to a host on an attached network. If the host is not on a directly attached network, an error is returned. This option can be used to send ping to a local host through an interface that has no route through it (for example, after the interface was dropped by `ROUTED`).

**-s packetsize**

Specifies the number of data bytes to be sent. The default is 56, which translates into 64 ICMP data bytes when combined with the 8 bytes of ICMP header data.

**-u**

Displays the time in microseconds (three decimal places). In order to ensure this microsecond precision, the `NTP_TIME` and `MICRO_TIME` kernel options must be on. By default, `NTP_TIME` and `MICRO_TIME` kernel options are off. If these

## ping

kernel options are off and this flag is used, the time is displayed to three decimal places, but in milliseconds.

**-v**

Specifies detailed output. ICMP packets other than ECHO\_RESPONSE that are received are listed.

## Examples

The following example shows how to use the ping command.

1. TCPIP> ping  
PING rufus.lkg.dec.com (10.10.2.4): 56 data bytes  
64 bytes from 10.10.2.4: icmp\_seq=0 ttl=64 time=30 ms  
64 bytes from 10.10.2.4: icmp\_seq=1 ttl=64 time=0 ms  
64 bytes from 10.10.2.4: icmp\_seq=2 ttl=64 time=0 ms  
64 bytes from 10.10.2.4: icmp\_seq=3 ttl=64 time=0 ms  
  
----rufus.lkg.dec.com PING Statistics----  
4 packets transmitted, 4 packets received, 0% packet loss  
round-trip (ms) min/avg/max = 0/8/30 ms

---

## route

Manually manipulates the routing tables.

### Format

#### Adding a Route

```
route [-nqvC] add [-net|-host] [family] destination[/bitmask] gateway [-link] [modifiers [args]]
```

#### Changing a Route

```
route [-nqv] change [-net|-host] [family] destination gateway [-link] [modifiers args]
```

#### Monitoring a Route

```
route [-n] monitor
```

#### Deleting a Route

```
route [-nqvC] delete [-net|-host] [family] destination[/bitmask] [-link] gateway [modifiers args]
```

#### Deleting All Routes

```
route [-nqvC] flush [family]
```

### Description

The `route` command allows you to manually manipulate the network routing tables. It normally is not needed, since a system-routing table management daemon, such as GATED or ROUTED, should tend to this task.

The `route` command accepts five commands:

|                      |                                                                                                         |
|----------------------|---------------------------------------------------------------------------------------------------------|
| <code>add</code>     | Adds a route.                                                                                           |
| <code>flush</code>   | Removes all routes.                                                                                     |
| <code>delete</code>  | Deletes a specific route.                                                                               |
| <code>change</code>  | Changes aspects of a route (such as its gateway).                                                       |
| <code>monitor</code> | Reports any changes to the routing information base, r' lookup misses, or suspected network partitions. |

The `flush` command has the following format:

```
route [-n] flush [family]
```

In this format, the address family can be specified by the `-inet` keyword.

The other commands have the following format:

```
route [-n] command [-net | -host] destination[/bitmask] gateway modifier [-netmask mask]
```

Unless the `net` or the `host` parameter is specified on the command line, `route` creates a host route or a network route by interpreting the Internet address associated with `destination` parameter. If the destination has a local address part of `INADDR_ANY`, or if the destination is the symbolic name of a network, a network route is created; otherwise, a host route is created.

For example, `128.32` is interpreted as `-host 128.0.0.32`, `128.32.130` is interpreted as `-host 128.32.0.130`; `-net 128.32` is interpreted as `128.32.0.0`, and `-net 128.32.130` is interpreted as `128.32.130.0`.

## route

All symbolic names specified for a destination or gateway are looked up first as a host name using `gethostbyname( )`. If this lookup fails, `getnetbyname( )` is then used to interpret the name as that of a network.

Routes added with the `route` command are marked as `RTF_STATIC` to differentiate them from routes added by the routing daemons (`GATED` or `ROUTED`). The `GATED` daemon does not remove the `RTF_STATIC` routes when it is shut down.

The `route` utility uses a routing socket and the new message types `RTM_ADD`, `RTM_DELETE`, and `RTM_CHANGE`. As such, only a privileged user can modify the routing tables.

## Flags

**-n**

Prevents attempts to display host and network names symbolically when reporting actions.

**-v**

Displays additional details.

**-q**

Suppresses all output.

## Modifiers

**-all**

Specifies that the kernel add or delete the specified route on all interfaces (for example, `TU0` and `TU1`) that are in the same subnet as the gateway. Use this modifier only with the `add` and `delete` modifiers. Do not use `-all` with the `-dev` and `-olddev` modifiers.

**-blackhole**

Specifies that this route is a blackhole route. Packets sent to blackhole routes are dropped, and notification is not sent to the packet originators. This is different from a normal route, which allows packets to be forwarded out on it. You must specify `127.0.0.1` (localhost) as the gateway argument.

**-cloning**

Generates a new route on use of this route.

**-dev *device***

Specifies the interface device to use in the routing entry. Use this modifier when you want to designate a particular interface for a route. If you do not specify this modifier, the route is added on the first interface that is found.

**-genmask *mask***

Specifies that the netmask *mask* is used for all routes cloned from this route.

**-hopcount *count***

Sets this route's maximum hopcount to the value specified by *count*.

**-iface | -interface**

Specifies that this route is through an interface instead of through a gateway (gateway is the default). This means the destination is reachable directly through

an interface; no intermediate system is required. The gateway parameter is the host address to be used for transmission.

**-inet**

Sets this route's type as AF\_INET. When you specify `inet` with the `delete` or `flush` command, only AF\_INET routes are deleted.

**-llinfo**

Specifies that this route contains valid link-layer information.

**-lock**

Locks the metric specified by the next modifier specified on the command line. A locked metric is not modified by the kernel. The following metrics can be locked: `mtu`, `hopcount`, `recvpipe`, `sendpipe`, `ssthresh`, `rtt`, and `rttvar`.

**-lockrest**

Locks the metrics specified by all the modifiers that follow on the command line. A locked metric is not modified by the kernel. The following metrics can be locked: `mtu`, `hopcount`, `recvpipe`, `sendpipe`, `ssthresh`, `rtt`, and `rttvar`.

**-mtu *size***

Sets this route's maximum transmission unit (MTU) (in bytes) to the value specified by *size*.

**-netmask *mask***

Specifies the subnet mask to use for the routing entry. Networks that use a nonstandard subnet must include this modifier. Specify this modifier after any optional modifiers. Do not specify this modifier if you specify a CIDR bitmask (*bitmask*). Do not specify this modifier with the `change` command.

**-nofragtopmtu**

Specifies that IP datagram fragmentation is disabled for this route.

**-nopmtudisc**

Specifies that path MTU discovery is disabled for this route.

**-olddev *device***

Specifies the old interface device that you want to change in the routing entry. Use this modifier with the `change` command only to move a route from one interface to another.

**-oldgateway *name***

Specifies the old gateway that you want to change in the routing entry. Use this modifier with the `change` command only.

**-precedence *value***

Sets the precedence of the route to the value specified by *value*. Among equivalent routes to the same destination, the route with the lower precedence is preferred.

**-recvpipe *bandwidth***

Sets this route's inbound delay bandwidth product (in bytes) to the value specified by *bandwidth*.

## route

### **-reject**

Specifies that this route is a reject route. Packets sent to reject routes are dropped, and messages designating the route as unreachable are sent to the packet originators. This is different from a normal route, which allows packets to be forwarded out on it. You must specify 127.0.0.1 (localhost) as the gateway argument.

### **-rtt *time***

Sets this route's round-trip time (in microseconds) to the value specified by *time*.

### **-rttvar *variance***

Sets this route's round-trip time variance (in microseconds) to the value specified by *variance*.

### **-sendpipe *bandwidth***

Sets this route's outbound delay bandwidth product (in bytes) to the value specified by *bandwidth*.

### **-ssthresh *threshold***

Sets this route's outbound gateway buffer limit (in bytes) to the value specified by *threshold*.

## Examples

The following examples show how to use the route utility.

1. TCPIP> route add default 128.32.0.130

The example shows how to add gateway 128.32.0.130 as a default gateway.

2. TCPIP> route add -host milan 128.32.0.130

The example shows how to add a route to host milan via gateway 128.32.0.130.

3. TCPIP> route delete -host milan 128.32.0.130

The example shows how to delete an existing route via gateway 128.32.0.130 to host milan.

4. TCPIP> route add -precedence 1 -host milan 128.32.0.130

The example shows how to add a route with a precedence value of 1 to host milan via gateway 128.32.0.130.

5. TCPIP> route change -oldgateway 128.32.0.130 -oldinterface le0 -  
\_TCPIP> -host milan 128.32.10.101

The example shows how to change an existing route for host milan via gateway 128.32.0.130 to use a new gateway 128.32.10.101.

---

## sysconfig

Maintains the kernel subsystem configuration.

### Format

```
sysconfig -c | -q | -Q | -r | -s | -u [subsystem-name] [attribute-list]
```

### Description

The `sysconfig` command is used to query and modify the kernel subsystem configuration. You use this command to add subsystems to your running kernel, reconfigure subsystems already in the kernel, ask for information about (query) subsystems in the kernel, and unconfigure and remove subsystems from the kernel.

A subset of kernel subsystems can be managed using the `sysconfig` command. This command allows you to modify the value of subsystem attributes, as long as the subsystem supports run-time modifications.

The first argument to the `sysconfig` command is the `subsystem-name` argument. The `subsystem-name` argument names the subsystem on which you want to perform the operation specified by one of the required flags, such as the `-c` (configure) flag or the `-q` (query attributes) flag. The `subsystem-name` argument is required for all flags except `-s` and `-m`. If you omit `subsystem-name` when you use one of these flags, the `sysconfig` command displays information about all loaded subsystems.

The `attribute-list` argument lists attribute names and, depending on the operation, attribute values. For reconfigure (`-r`) operations, the `attribute-list` argument has the following format:

```
attribute1=value1 attribute2=value2...
```

You cannot include spaces between the attribute name, the equals sign (=), and the value.

For query attribute (`-q`) operations, the `attribute-list` argument has the following format:

```
attribute1 attribute2...
```

The `attribute-list` argument is required when you use the `-r` flag and is optional with the `-q` flag. Any `attribute-list` specified with other flags is ignored by the `sysconfig` command.

When you configure a subsystem using the `-c` flag, you make that subsystem available for use. If the subsystem is loadable, the `sysconfig` command loads the subsystem and then initializes the value of its attributes.

If you want to modify the value of a subsystem attribute, you use the `-r` (reconfigure) flag. When you use the `-r` flag, the `sysconfig` command modifies the named attributes by storing the value you specify in them. The modifications take effect immediately.

To get information about subsystem attributes, use either the `-q` flag or the `-Q` flag. You can specify an attribute list with both these flags. When you use the `-q` flag, the `sysconfig` command reads the value of attributes from the kernel and displays those values on your local display. When you use the `-Q` flag,

## sysconfig

the `sysconfig` command displays the following information about either each attribute in the subsystem or, if specified, each attribute in the `attribute-list`:

- Attribute datatype.
- Operations supported by the attribute. This information indicates, for example, whether you can reconfigure the attribute using the `sysconfig -r` command.
- Minimum and maximum allowed attribute value.

To get information about the state of subsystems, use the `-s` flag. This flag provides a list of the subsystems that are currently loaded and configured into the kernel. If you specify `subsystem-name`, the command displays information about the state of that subsystem. Each subsystem can have one of three states:

- Loaded and configured (available for use)
- Loaded and unconfigured (not available for use but still loaded)  
This state applies only to static subsystems, which can be unconfigured but cannot be unloaded.
- Unloaded (not available for use)  
This state applies only to loadable subsystems, which are automatically unloaded from the kernel when you unconfigure them with the `sysconfig -u` command.

Subsystems that are not being used can be unconfigured using the `-u` flag. Unconfiguring subsystems can help save kernel memory, making it available for other uses. You can unconfigure any static or loadable subsystem that supports run-time unconfiguration. If you unconfigure a loadable subsystem, that subsystem is also unloaded from the kernel.

You can use the `sysconfig` command to display the value of attributes on the local system. If you want to configure, reconfigure, or unconfigure a subsystem, you must be authorized to modify the kernel configuration. Only users who have a system group UIC or who have an account with `SYSPRV`, `BYPASS`, `OPER` privileges can configure, reconfigure, or unconfigure the subsystems.

## Flags

### **-c**

Configures the specified subsystem by initializing its attribute values and, possibly, loading it into the kernel.

### **-q**

Queries attribute values for the configured subsystem specified by `subsystem-name`. If you omit `attribute-list`, values for all the specified subsystem's attributes are displayed.

### **-Q**

Queries information about attributes of the configured subsystem specified by `subsystem-name`. The information includes the attribute data type, the operations supported, and the minimum and maximum values allowed for the attribute. Note that the minimum and maximum values refer to length and size for attributes of char and binary types, respectively. If you omit the `attribute-list` argument, information about all attributes in the specified subsystem is displayed.

**-r**

Reconfigures the specified subsystem. You must supply the `subsystem-name` argument and the `attribute-list` argument when you use this flag.

**-s**

Queries the subsystem state for the specified subsystems. If you omit the subsystem name, `sysconfig` displays the state of all the configured subsystems.

**-u**

Unconfigures and, if the subsystem is loadable, unloads the specified subsystem from the kernel.

## Examples

The following examples show how to use the `sysconfig` command.

1. 

```
TCPIP> sysconfig -s
inet: loaded and configured
net: loaded and configured
socket: loaded and configured
iptunnel: loaded and configured
ipv6: loaded and configured
snmpinfo: loaded and configured
```

This example shows how to display the kernel subsystems and their status.

2. 

```
TCPIP> sysconfig -q net
net:
ifnet_debug = 0
ifqmaxlen = 1024
lo_devs = 1
lo_def_ip_mtu = 4096
nslip = 0
TCPIP>
```

This example shows how to display subsystem attributes and their values.

3. 

```
TCPIP> sysconfig -s net
net: loaded and configured
%SYSTEM-F-NOPRIV, insufficient privilege or object protection violation
TCPIP>
```

This example shows how to query the state of a particular subsystem.

4. 

```
TCPIP> sysconfig -u net
Unable to find subsystem NET
%SYSTEM-F-NOPRIV, insufficient privilege or object protection violation
TCPIP>
```

This example shows how to unconfigure a particular subsystem. You must have system privileges to display the output from this command.

## TCPTRACE

---

### TCPTRACE

Traces packets between two hosts.

#### Format

```
TCPTRACE host [/BUFFERS=n | /FULL | /OUTPUT=file | /PACKETS=n | /PORT=option |
/PROTOCOL=option]
```

#### Description

TCPTRACE traces packets as they travel between the local and remote host. You can trace all packets or you can use command qualifiers to monitor only those packets of interest.

#### Qualifiers

##### **/BUFFERS=*n***

Optional. The default is 100.

Specifies the number of buffers that TCPTRACE allocates for temporary storage.

These buffers must be locked into the working set, so the number can be:

- Decreased (to be accommodated in the working set)
- Raised (to prevent the dropping of trace packets)

##### **/FULL**

Optional. The default is brief display.

Displays the packet's contents.

##### **/OUTPUT=*file***

Optional. The default is screen display. Redirects the output from screen to the specified file. If this file name already exists, the output is appended to it.

##### **/PACKETS=*n***

Optional. The default is 10.

Stops the trace after the specified number of packets is displayed.

##### **/PORT=*option***

{LOCAL=*n* | REMOTE=*n*}

Optional for port number. The default is that all traffic is displayed.

Required for port type. Filters the trace to the specified port.

##### **/PROTOCOL=*option***

{ARP | ICMP | IP | TCP | UDP}

Optional. The default is /PROTOCOL=IP.

Filters the specified protocol.

## Examples

The following examples show how to use the TCPTRACE command.

1. `$ TCPTRACE HOST1 /FULL /PORT=(REMOTE=21)`

This example shows how to use the TCPTRACE command to trace packets between the local system and Host1. TCPTRACE filters all packets except those packets directed to port 21 on the remote host.

2. `$ TCPTRACE HOST2 /PORT=(LOCAL=23, REMOTE=1056) -  
_ $ /FULL /PACKETS=30 /OUTPUT=TELNET_TRACE.TXT`

This example shows how to use the TCPTRACE command to trace packets between the local system and Host2. TCPTRACE filters all packets except those packets directed to port 23 on the local host and port 1056 on the remote host. The trace results are output to a file with the name TELNET\_TRACE.TXT. The trace stops after 30 packets meeting the specifications are encountered.

## traceroute

---

### traceroute

Displays the route that packets take to the network host.

#### Format

```
traceroute [-m max_ttl] [-n] [-p port] [-q nqueries] [-r] [-s src_addr] [-v] [-w waittime] host [packetsize]
```

#### Description

The Internet is a large and complex aggregation of network hardware connected together by gateways. The `traceroute` command tracks the route that packets follow from gateway to gateway. The command uses the IP protocol time-to-live (TTL) field and attempts to elicit an ICMP `TIME_EXCEEDED` response from each gateway along the path to a particular host.

The only mandatory parameter is the destination host name or IP address. The default probe datagram length is 38 bytes, but you can increase this value by specifying a packet size (in bytes) after the destination host name. This is useful when the `-f` option is given for MTU discovery along the route. You should start with the maximum packet size for your own network interface (if the given value is even bigger, `traceroute` attempts to select a more appropriate value). If no packet size is given when using the `-f` option, `traceroute` determines the initial MTU automatically.

To track the route of an IP packet, `traceroute` launches UDP probe packets with a small TTL (time-to-live) and then listens for an ICMP “time exceeded” reply from a gateway. Probes start with a TTL of 1 and increment by one until either an ICMP “port unreachable” is returned (indicating that the packet reached the host) or until the maximum number of hops is exceeded (the default is 30 hops and can be changed with the `-m` option). At each TTL setting, three probes are launched (the number can be changed with the `-q` option), and `traceroute` displays a line showing the TTL, address of the gateway, and round-trip time of each probe. If the probe answers come from different gateways, `traceroute` displays the address of each responding system. If there is no response within a 3-second timeout interval (which can be changed with the `-w` option), `traceroute` displays an asterisk (\*) for that probe.

To prevent the destination host from processing the UDP probe packets, the destination port is set to an unlikely value. If necessary, you can change the destination port valued with the `-p` option.

---

#### Note

---

This program is intended for use in network testing, measurement, and management. It should be used primarily for manual fault isolation. Because of the load it could impose on the network, do not use `traceroute` during normal operations or from automated scripts.

---

#### Flags

##### **-A**

Looks up the AS-number (autonomous system) for each hop’s network address at the `whois` server specified by the `-h` option.

**-a**

If the destination host has multiple addresses, `tracert` probes all addresses if this option is set. Normally, only the first address as returned by the resolver is attempted.

**-c stoptime**

Specifies a delay (in seconds) to pause between probe packets. This can be necessary if the final destination is a router that does not accept undeliverable packets in bursts.

**-f**

Disables IP fragmentation. If the given packet size is too big to be handled unfragmented by a machine along the route, a "fragmentation needed" status is returned, and the indicator !F is printed. If a gateway returns the proper MTU size to be used, `tracert` automatically decreases the packet size to this new value. If the proper MTU size is not returned, `tracert` chooses a smaller packet size.

**-g gateway**

Enables the IP LSRR (loose source record route) option. This is useful for asking how somebody at the specified gateway reaches a particular target.

**-h server**

Specifies the name or IP address of the `whois` server that is contacted for the AS-number lookup, if the `-A` option is given.

**-i initial\_ttl**

Sets the starting time-to-live value to `initial_ttl`, to override the default value of 1. Effectively this skips processing for intermediate hosts that are less than `initial_ttl` hops away.

**-k**

Keeps the connection to the `whois` server permanently open. This speeds lookups considerably, because a connection setup for each individual lookup is not necessary. However, not all `whois` servers support this feature.

**-l**

Prints the value of the TTL field in each packet received. (This flag can be used to help detect asymmetric routing.)

**-m max\_ttl**

Sets the maximum time-to-live (maximum number of hops) used in outgoing probe packets. The default is 30 hops, which is the same default used for TCP connections.

**-N**

Displays the network name for each hop. If a BIND resolver cannot be reached, network names are retrieved just from the `/etc/networks` file.

**-n**

Displays the hop IP addresses using dotted-decimal notation. This saves a name server address-to-name lookup for each gateway found on the path. It also prevents a reverse lookup for numeric dotted-quad addresses given on the command line, such as destination host or `-g` gateway addresses.

## traceroute

### **-p port**

Sets the base UDP port number used in probes. (The default value is 33434.) The `traceroute` command presumes that nothing is listening on UDP ports `base` to `base+n hops-1` at the destination host (so an ICMP “port unreachable” message is returned to terminate the route tracing). If another process is listening on a port in the default range, use this option to pick an unused port range.

### **-Q maxquit**

Stops probing this hop after the number of consecutive timeouts specified by `maxquit` are detected. The default value is 5. Useful in combination with `-S` if you have specified a big `nqueries` probe count.

### **-q nqueries**

Sets the number of probes launched at each TTL setting. The default is 3.

### **-r**

Bypasses the normal routing tables and sends directly to a host on an attached network. If the host is not on a directly attached network, an error is returned. This option can be used to ping a local host through an interface that has no route through it (for example, after the interface was dropped by `ROUTED` or `GATED`).

### **-S**

Prints a summary of per-hop minimum/average/maximum rtt (round-trip time) statistics. This flag suppresses the per-probe rtt and TTL reporting. To obtain more detailed statistics, increase the default `nqueries` probe count. For more information, see also the `-Q` option.

### **-s src\_addr**

Uses the following IP address (which must be given as an IP number, not as a host name) as the source address in outgoing probe packets. On hosts with more than one IP address, this option can be used to force the source address to be something other than the IP address of the interface on which the probe packet is sent. If the IP address is not one of this machine’s interface addresses, an error is returned and nothing is sent.

### **-v**

Lists any received ICMP packets other than `TIME_EXCEEDED` and `UNREACHABLE`.

### **-w waittime**

Sets the time (in seconds) to wait for a response to a probe. The default is 3 seconds.

## Parameters

### **host**

Specifies the name or IP address of the destination host. This parameter is required.

### **packetsize**

Specifies the default length for a probe datagram. This parameter is optional. The default is 38 bytes.

## Examples

The following examples show how to use the tracert command.

```
1. localhost> tracert nis.nsf.net
tracert to nis.nsf.net (35.1.1.48), 30 hops max, 56 byte packet
 1 helios.ee.lbl.gov (128.3.112.1) 19 ms 19 ms 0 ms
 2 lilac-dmc.Berkeley.EDU (128.32.216.1) 39 ms 39 ms 19 ms
 3 lilac-dmc.Berkeley.EDU (128.32.216.1) 39 ms 39 ms 19 ms
 4 ccngw-ner-cc.Berkeley.EDU (128.32.136.23) 39 ms 40 ms 39 ms
 5 ccn-nerif22.Berkeley.EDU (128.32.168.22) 39 ms 39 ms 39 ms
 6 128.32.197.4 (128.32.197.4) 40 ms 59 ms 59 ms
 7 131.119.2.5 (131.119.2.5) 59 ms 59 ms 59 ms
 8 129.140.70.13 (129.140.70.13) 99 ms 99 ms 80 ms
 9 129.140.71.6 (129.140.71.6) 139 ms 239 ms 319 ms
10 129.140.81.7 (129.140.81.7) 220 ms 199 ms 199 ms
11 nic.merit.edu (35.1.1.48) 239 ms 239 ms 239 ms
```

This tracert command displays the route that packets take to a remote host. In this example, note that display lines 2 and 3 are identical. This is due to a bug in the kernel on the second hop system, lbl-csam.arpa, that forwards packets with a zero TTL. (This is a bug in the distributed version of BSD Version 4.3.)

```
2. localhost> tracert allspice.lcs.mit.edu
tracert to allspice.lcs.mit.edu (18.26.0.115), 30 hops max
 1 helios.ee.lbl.gov (128.3.112.1) 0 ms 0 ms 0 ms
 2 lilac-dmc.Berkeley.EDU (128.32.216.1) 19 ms 19 ms 19 ms
 3 lilac-dmc.Berkeley.EDU (128.32.216.1) 39 ms 19 ms 19 ms
 4 ccngw-ner-cc.Berkeley.EDU (128.32.136.23) 19 ms 39 ms 39 ms
 5 ccn-nerif22.Berkeley.EDU (128.32.168.22) 20 ms 39 ms 39 ms
 6 128.32.197.4 (128.32.197.4) 59 ms 119 ms 39 ms
 7 131.119.2.5 (131.119.2.5) 59 ms 59 ms 39 ms
 8 129.140.70.13 (129.140.70.13) 80 ms 79 ms 99 ms
 9 129.140.71.6 (129.140.71.6) 139 ms 139 ms 159 ms
10 129.140.81.7 (129.140.81.7) 199 ms 180 ms 300 ms
11 129.140.72.17 (129.140.72.17) 300 ms 239 ms 239 ms
12 * * *
13 128.121.54.72 (128.121.54.72) 259 ms 499 ms 279 ms
14 * * *
15 * * *
16 * * *
17 * * *
18 ALLSPICE.LCS.MIT.EDU (18.26.0.115) 339 ms 279 ms 279 ms
```

In this example, gateways 12, 14, 15, 16, and 17 either do not send ICMP “time exceeded” messages or send them with a TTL too small to reach the local host. Further investigation is required to determine the cause. For example, by contacting the system administrators for gateways 14 through 17, you could discover that these gateways are running the MIT C Gateway code that does not send “time exceeded” messages.



---

## Commonly Used UNIX Commands

Table B-1 contains some commonly used TCP/IP network management commands and their corresponding Compaq *Tru64* UNIX command formats. A UNIX network manager or system manager who is unfamiliar with the TCP/IP Services implementation of TCP/IP might find this information helpful.

Please note that TCP/IP management commands and qualifiers are a subset of the Compaq *Tru64* UNIX network commands. Therefore, not all Compaq *Tru64* UNIX command options have a corresponding TCP/IP management command. In addition, execution of a UNIX command in a DCL command procedure does not return an error in \$STATUS, so you cannot test for the failure of a UNIX command in a DCL command procedure.

TCP/IP Services uses pseudointerfaces to associate more than one IP address with an interface. The TCP/IP commands SET INTERFACE and SET NOINTERFACE allow you to add or delete one IP address at a time. Therefore, to add or delete multiple pseudointerfaces, you must enter a SET INTERFACE or SET NOINTERFACE command for each pseudointerface. For information about how to initialize an interface with more than one IP address, refer to the *DIGITAL TCP/IP Services for OpenVMS Concepts and Planning* and *Compaq TCP/IP Services for OpenVMS Management* manuals.

The Compaq *Tru64* UNIX `ifconfig` command allows you to associate multiple IP addresses per interface by entering one `ifconfig` command. Therefore, to create or delete multiple IP addresses in one command, use the `ifconfig` command.

## Commonly Used UNIX Commands

**Table B-1 Commonly Used Commands**

| TCP/IP Command                                                                                          | UNIX Command                                                                                        | Description                                             |
|---------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| <b>ARP information</b>                                                                                  |                                                                                                     |                                                         |
| SHOW ARP                                                                                                | netstat -Hn                                                                                         | Displays all ARP entries. <sup>1</sup>                  |
| SHOW ARP                                                                                                | arp -a[-u][-i][-n]                                                                                  | Displays all ARP entries.                               |
| SHOW ARP <i>hostname</i>                                                                                | arp <i>hostname</i>                                                                                 | Displays entry for a host. <sup>1</sup>                 |
| SET ARP <i>hardware_addr host</i>                                                                       | arp -s <i>host hardware_addr [params]</i>                                                           | Adds an ARP entry.                                      |
| SET NOARP <i>host</i>                                                                                   | arp -d <i>host</i>                                                                                  | Removes an ARP entry.                                   |
| <b>Interface information</b>                                                                            |                                                                                                     |                                                         |
| SHOW INTERFACE                                                                                          | ifconfig -a[-d][-u][-v]<br>[-f <i>addr_family</i> ]                                                 | Displays information for all interfaces.                |
| SHOW INTERFACE                                                                                          | ifconfig -l[-d][-u][-v]<br>[-f <i>addr_family</i> ]                                                 | Displays information for all configured interfaces.     |
| SHOW INTERFACE /FULL                                                                                    | netstat -i[-n][-d]                                                                                  | Displays information for all configured interfaces.     |
| SHOW INTERFACE /FULL                                                                                    | netstat -I <i>interface</i>                                                                         | Displays information on the specified interface.        |
|                                                                                                         | netstat -A                                                                                          | Displays miscellaneous information.                     |
|                                                                                                         | netstat -I <i>interface</i> -t                                                                      | Displays timer information for the specified interface. |
| SHOW INTERFACE <i>interface_id</i> [/FULL]                                                              | ifconfig <i>interface_id</i>                                                                        | Displays interface information.                         |
| SET INTERFACE <i>interface_id</i><br>/HOST= <i>host</i><br>/NETWORK_MASK= <i>mask</i>                   | ifconfig <i>interface_id</i><br>[ <i>address_family</i> ] <i>address</i> [/ <i>bitmask</i> ]        | Sets the primary interface address.                     |
| SET INTERFACE <i>interface_id</i><br>/HOST= <i>host</i><br>/[NETWORK_MASK= <i>mask</i> ]                | ifconfig <i>interface_id</i><br>[ <i>address_family</i> ] <i>alias</i> <i>addr</i> [/ <i>mask</i> ] | Adds an alias interface address.                        |
| SET NOINTERFACE <i>interface_id</i>                                                                     | ifconfig <i>interface_id</i><br>[ <i>address_family</i> ]<br>-alias <i>alias_address</i>            | Removes an alias interface address.                     |
| SET NOINTERFACE <i>interface_id</i><br>/HOST= <i>host</i><br>/[NETWORK_MASK= <i>mask</i> ] <sup>2</sup> | ifconfig <i>interface_id</i><br>[ <i>address_family</i> ]<br>-aliaslist <i>alias_list/bitmask</i>   | Adds a list of interface aliases.                       |
| SET NOINTERFACE <i>interface_id</i> <sup>3</sup>                                                        | ifconfig <i>interface_id</i><br>[ <i>address_family</i> ]<br>-aliaslist <i>alias_list/bitmask</i>   | Removes a list of interface alias addresses.            |
| SET INTERFACE <i>interface_id</i><br>/PARAM= <i>value</i>                                               | ifconfig <i>interface_id</i> <i>param=value</i>                                                     | Modifies interface parameters.                          |

<sup>1</sup>Displays the same information for both the Compaq *Tru64* UNIX and the TCP/IP versions of the command.

<sup>2</sup>To create a new pseudointerface for each alias, enter a command for each alias address.

<sup>3</sup>To delete a pseudointerface for each alias, enter a command for each alias.

(continued on next page)

**Table B–1 (Cont.) Commonly Used Commands**

| TCP/IP Command                                                                                  | UNIX Command                                                                                               | Description                                                        |
|-------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------|
| <b>Active sockets information</b>                                                               |                                                                                                            |                                                                    |
| SHOW DEVICE[/FULL]                                                                              | netstat -a                                                                                                 | Displays the active sockets.                                       |
| SHOW DEVICE[/FULL]                                                                              | netstat -n                                                                                                 | Displays the active sockets.                                       |
| SHOW DEVICE BG <sub>xx</sub> /[FULL]                                                            |                                                                                                            | Displays information on a particular socket.                       |
| <b>Memory information</b>                                                                       |                                                                                                            |                                                                    |
| SHOW COMMUNICATION/MEMORY                                                                       | netstat -m                                                                                                 | Displays memory usage. <sup>1</sup>                                |
| <b>Route information</b>                                                                        |                                                                                                            |                                                                    |
| SHOW ROUTE[/FULL]                                                                               | netstat -r                                                                                                 | Displays all routes.                                               |
| SHOW ROUTE <i>host</i>                                                                          |                                                                                                            | Displays the route to a destination.                               |
| SHOW ROUTE/GATEWAY= <i>host</i>                                                                 |                                                                                                            | Displays routes.                                                   |
| SHOW ROUTE[/FULL]                                                                               | netstat -r -f <i>address_family</i>                                                                        | Displays the route for an address family.                          |
|                                                                                                 | netstat -rs                                                                                                | Displays route statistics.                                         |
|                                                                                                 | netstat -rv                                                                                                | Displays routes with metrics.                                      |
| SET ROUTE <i>destination</i><br>[/NETWORK]<br>[/MASK= <i>bitmask</i> ]<br>/GATEWAY= <i>host</i> | route add -net -host<br><i>destination_address/bitmask</i><br><i>gateway_addr</i>                          | Adds a route to a host.                                            |
| SET NOROUTE <i>destination</i><br>[/GATEWAY= <i>host</i> ]                                      | route delete -net -host<br><i>destination_address/bitmask</i><br>[ <i>gateway_addr</i> ]                   | Removes a route to a host.                                         |
|                                                                                                 | route [-nqv] change [-net -host]<br>[ <i>address_family</i> ] <i>destination</i><br><i>gateway</i> [-link] | Changes route information.                                         |
| <b>Protocol information</b>                                                                     |                                                                                                            |                                                                    |
| SHOW PROTOCOL                                                                                   | netstat -s                                                                                                 | Displays statistics for all protocols. <sup>1</sup>                |
| SHOW PROTOCOL <i>protocol_id</i>                                                                | netstat -p <i>protocol_id</i>                                                                              | Displays protocol statistics for a specific protocol. <sup>1</sup> |
| ZERO PROTOCOL <i>protocol_id</i>                                                                | netstat -p <i>protocol_id</i> -z                                                                           | Zeroes protocol statistics.                                        |
|                                                                                                 | sysconfig -s                                                                                               | Shows the configurable subsystems.                                 |

<sup>1</sup>Displays the same information for both the Compaq *Tru64* UNIX and the TCP/IP versions of the command.

(continued on next page)

## Commonly Used UNIX Commands

**Table B–1 (Cont.) Commonly Used Commands**

| TCP/IP Command                                                                                                                  | UNIX Command                                 | Description                                                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Protocol information</b>                                                                                                     |                                              |                                                                                                                          |
| SHOW PROTOCOL IP/PARAMETERS<br>SHOW PROTOCOL UDP/PARAMETERS<br>SHOW PROTOCOL TCP/PARAMETERS<br>SHOW PROTOCOL<br>ICMP/PARAMETERS | sysconfig -q inet [ <i>attribute-list</i> ]  | Shows the inet subsystem parameters.                                                                                     |
|                                                                                                                                 | sysconfig -c <i>attr=value ...</i>           | Configures subsystems.                                                                                                   |
|                                                                                                                                 | sysconfig -q <i>subsystem</i>                | Shows information about a subsystem.                                                                                     |
| SET PROTOCOL <i>protocol_id</i><br>/PARAMETERS= <i>value</i>                                                                    | sysconfig -r inet <i>attr=value ...</i>      | Sets protocol attributes.                                                                                                |
|                                                                                                                                 | sysconfig -r <i>subsystem attr=value ...</i> | Sets other subsystem attributes.                                                                                         |
| <b>NFS information</b>                                                                                                          |                                              |                                                                                                                          |
| SET NFS_SERVER<br>/UID_DEFAULT= <i>value</i>                                                                                    | sysconfig -r nfs noproxy_uid= <i>value</i>   | Specifies the default UID when a user cannot be translated by the proxy.                                                 |
| SET NFS_SERVER<br>/GID_DEFAULT= <i>value</i>                                                                                    | sysconfig -r nfs noproxy_gid= <i>value</i>   | Specifies the default GID when a user cannot be translated by the proxy.                                                 |
| SET NFS_SERVER<br>/THREADS=TCP= <i>value</i>                                                                                    | sysconfig -r nfs tcp_threads= <i>value</i>   | Specifies the number of concurrent TCP threads within the NFS server. A value of 0 (zero) will disable the TCP protocol. |
| SET NFS_SERVER<br>/INACTIVITY_TIMER= <i>value</i>                                                                               | sysconfig -r nfs vnode_age= <i>value</i>     | Specifies the time interval (in seconds) since the last file access request.                                             |

## A

---

arp command, 1-2, 1-6, A-2

### Attribute

- displaying values, 2-2
- modifying temporarily, 2-2
- testing support, 2-2

### Auxiliary server

- increasing connection limit, 2-14

## B

---

BIND log file, 1-9

## C

---

Checking interface parameters, 1-5

### Commands

- commonly used, B-2

### Connection limit

- increasing, 2-14

### Connection ports

- outgoing
  - modifying range, 2-10

### Connection Timeout Rate

- increasing the TCP partial, 2-7

### Context timeout rate

- increasing, 2-9

## D

---

Determining if TCP/IP is running, 1-1

### Diagnostic tools, 1-2

- arp command, 1-2, 1-6, A-2
- dig command, 1-2, A-5
- ifconfig command, 1-2, A-11
- ndc command, 1-2, A-16
- netstat command, 1-2, A-18
- nslookup command, 1-2, A-22
- ping command, 1-2, 1-3, A-27
- route command, 1-2, A-31
- sysconfig command, 1-2, A-35
- TCPTRACE command, 1-2, A-38
- traceroute command, 1-2, A-40
- dig command, 1-2, A-5

### Displaying

- allocated sockets, 1-7
- attribute support, 2-2
- attribute values, 2-2
- hardware addresses, 1-6
- IP addresses, 1-6
- network statistics, 1-6, 2-11
- routes, 1-10
- socket statistics, 2-13
- subsystems and attributes, 2-1
- the arp table, 1-6

## E

---

### Errors

- BIND log file, 1-9

## I

---

ifconfig command, 1-2, 1-5, A-11

### Improving performance

- benefits and tradeoffs, 2-3

inet subsystem, 2-1, 2-6, 2-15

- attributes, 2-6

inifaddr\_hsize attribute, 2-6

Internet subsystem, 2-6

### IP addresses

- logging, 2-13

ipport\_userreserved attribute, 2-6, 2-10, 2-15

ipport\_userreserved\_min attribute, 2-6, 2-10

ipqmaxlen attribute, 2-6

ipqs attribute, 2-6

iptunnel subsystem, 2-1

ipv6 subsystem, 2-1

## K

---

### Keepalive function

- enabling, 2-9

Kernel attributes, 2-1

Kernel interface alias table

- increasing the size, 2-7

Kernel subsystems

- modifying, 2-3

## M

---

### Modifying

- attribute values, 2-3
- internet subsystem attributes, 2-6
- kernel subsystems, 2-3
- outgoing connection ports, 2-10
- socket subsystem attributes, 2-4

### Monitoring

- servers, 2-11

### Multihomed host

- testing interfaces, 1-5

## N

---

### Name server

- log file, 1-9

### ndc command, 1-2, A-16

### netstat command, 1-2, 1-6, 1-10, 2-11, 2-13, A-18

### net subsystem, 2-1

### Network management tasks

- commonly performed, B-2

### Network statistics

- displaying, 2-11

### Network tuning guidelines, 2-3

### nslookup command, 1-2, A-22

### NSLOOKUP utility

- set commands, A-24

## P

---

### Pending TCP connections, 2-5

### Performance

- benefits and tradeoffs, 2-3

### ping command, 1-2, 1-3, 1-8, A-27

### PMTU protocol, 2-10

### pmtu\_enabled attribute, 2-6, 2-10, 2-15

### Problems

- application, 1-1
- configuration, 1-1
- device driver, 1-7
- isolating, 1-1
- loss of connectivity, 1-3
- memory usage, 1-7
- misconfigured interface, 1-5
- name server, 1-8
- network collision errors, 1-7
- network failure, 1-1
- network input/output errors, 1-7
- network service not responding, 1-11
- network unreachable, 1-10
- physical failure, 1-1
- solving performance, 2-14
- transport failure, 1-1
- unavailable network service, 1-11
- user error, 1-1

## R

---

### Retransmission rate

- decreasing, 2-8

### route command, 1-2, A-31

### Routing table

- displaying routes, 1-10

## S

---

### sbcompress\_threshold attribute, 2-4, 2-15

### sb\_max attribute, 2-4, 2-5

### Server applications

- configuring memory, 2-13
- logging IP addresses, 2-13
- tuning, 2-13

### Servers, 2-11

### Service configuration parameters, 1-12

### Services database, 1-12

### SHOW ARP command, 1-6

### SHOW DEVICE command, 1-1

### SHOW INTERFACE command, 1-6

### SHOW NAME\_SERVICE command, 1-8

### SHOW ROUTE command, 1-10

### SHOW SERVICE /FULL command, 1-13

### SHOW SERVICE command, 1-11

### snmpinfo subsystem, 2-1

### sobacklog\_drops attribute, 2-11, 2-13

### sobacklog\_hiwat attribute, 2-11, 2-13

### Socket

- displaying statistics, 2-13

### Socket buffer

- increasing maximum size, 2-5

### socket subsystem, 2-1, 2-4, 2-15

### somaxconn attribute, 2-4, 2-5, 2-7, 2-13, 2-15

### somaxconn\_drops attribute, 2-11, 2-13

### sominconn attribute, 2-4, 2-5, 2-13, 2-15

### sysconfig command, 1-2, A-35

## T

---

### tcbhadhnum attribute, 2-6

### tcbhashsize attribute, 2-6

### TCP hash table

- increasing number, 2-6
- increasing size, 2-6

### TCPTIME command, 1-2, 1-7, A-38

### tcp\_keepalive\_default attribute, 2-6, 2-9

### tcp\_keepcnt attribute, 2-9

### tcp\_keepidle attribute, 2-9

### tcp\_keepinit attribute, 2-6, 2-7, 2-9

### tcp\_keepintvl attribute, 2-9

### tcp\_msl attribute, 2-6

### tcp\_rexmit\_interval\_min attribute, 2-6, 2-8

### Testing connectivity, 1-3

Timeout rate  
  increasing the, 2-7  
traceroute command, 1-2, A-40  
Tracing packets, 1-7  
Troubleshooting process, 1-1  
Tuning  
  auxiliary server, 2-14

  guidelines, 2-3  
  server applications, 2-13  
Tuning recommendations  
  primary server, 2-15

## **U**

---

UNIX commands, B-2

