

TCP/IP Troubleshooting

In the 1990s, Microsoft set out to significantly improve the scalability of Microsoft networking by introducing a completely rewritten TCP/IP stack. The new TCP/IP stack, designed to incorporate many advances in performance and ease of administration, was a high-performance implementation of the industry-standard TCP/IP protocol.

With each generation of Microsoft® Windows®, Microsoft's implementation of the TCP/IP protocol stack has continued to evolve and include new features and services that enhance performance, security, and reliability. The TCP/IP protocol stack for the Microsoft® Windows Server™ 2003 operating system is self-tuning, more scalable, easier to administer, faster, and more secure. The TCP/IP protocol stack and its associated services are installed by default, and they can no longer be uninstalled by using the Network Connections feature.

As with previous versions of Windows server operating systems, a variety of diagnostic and repair tools are included to help you quickly isolate and resolve TCP/IP communication problems. In addition to the tools included in previous versions of Windows server, new tools and features have been added to assist you in troubleshooting TCP/IP communications.

This chapter discusses a variety of troubleshooting utilities and tools included in Windows Server 2003 and provides a basic structure that you can use to troubleshoot TCP/IP communication problems.

In This Chapter

TCP/IP Communication Process Overview	3
Troubleshooting Overview	10
Unable to Reach a Host or NetBIOS Name	28
Troubleshooting IP Routing.....	34
Troubleshooting Services	40
Additional Resources.....	41

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place or event is intended or should be inferred.

© 2004 Microsoft Corporation. All rights reserved.

Microsoft, Windows, Windows NT, Windows Server, and the Windows logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

TCP/IP Communication Process Overview

TCP/IP uses a consistent sequence to establish a communication link over a network or across disparate networks. Before sending out the first packet that will establish a communication session, the TCP/IP protocol on the sending host performs these four distinct steps:

1. TCP/IP resolves the host name or NetBIOS name to an IP address.
2. Using the destination IP address and the IP routing table, TCP/IP determines the interface to use and the next-hop IP address.
3. For unicast IP traffic on shared access technologies such as Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI), the Address Resolution Protocol (ARP) resolves the next-hop IP address to a media access control (MAC) address (also known as a data link-layer address).

For multicast IP traffic on Ethernet and FDDI, the destination multicast IP address is mapped to the appropriate multicast MAC address. For multicast IP traffic on Token Ring, the functional address of 0xC0-00-00-04-00-00 is used. For broadcast traffic on shared access technologies, the MAC address is mapped to 0xFF-FF-FF-FF-FF-FF.

4. The IP datagram is then sent to the MAC address resolved through ARP, to the multicast mapping, or to the MAC-level broadcast address.



Note

For an in-depth technical guide to TCP/IP protocols and services and their implementation in Windows Server 2003, see *Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference*, published by Microsoft Press.

TCP/IP Communication Process

The TCP/IP stack always follows the sequence described below to determine how to get a packet from point to point. If you are interested in the standard troubleshooting sequence, see “Unable to Reach a Host or NetBIOS Name” later in this chapter.

Resolving a Name to an IP Address

If the destination to be reached by a program is in the format of a NetBIOS name or host name, name resolution is required before IP can send the first packet. IP only understands IP addresses; host and NetBIOS names are each resolved to an IP address in different ways.

Resolving a NetBIOS Name to an IP Address

NetBIOS names can be directly resolved to an IP address through four mechanisms: consulting the NetBIOS name cache, querying a WINS server, broadcasting, or checking the Lmhosts file.

Windows Server 2003–based computers always begin by checking the host computer’s internal NetBIOS name cache. If this fails to provide an IP address, the NetBIOS name can be resolved to an IP address by using a WINS server, a series of broadcasts, or the Lmhosts file. Which of these three is used first by any particular computer depends on its node type. The default node type is hybrid or H-node, which starts by querying a WINS server, then attempts a local broadcast to resolve the name. If these methods do not work, a client converts the NetBIOS name to a host name and performs host-name resolution.

Computers running Windows Server 2003 are B-node by default. They become H-node when they are configured with a WINS server. The following node types are available:

- **B-node (broadcast).** B-node uses broadcast NetBIOS name queries for name registration and resolution. B-node has two major problems: Broadcasts disturb every node on the network; and routers typically do not forward broadcasts, so only NetBIOS names on the local network can be resolved.
- **P-node (peer-peer).** P-node uses a NetBIOS name server (NBNS), such as a WINS server, to resolve NetBIOS names. P-node does not use broadcasts; instead, it queries an NBNS directly.
- **M-node (mixed).** M-node is a combination of B-node and P-node. By default, an M-node functions as a B-node. If an M-node is unable to resolve a name by broadcast, it queries an NBNS by using P-node.
- **H-node (hybrid).** H-node is a combination of P-node and B-node. By default, an H-node functions as a P-node. If an H-node is unable to resolve a name through an NBNS, it uses a broadcast to resolve the name.

If the only problem is NetBIOS name resolution, the computer should still be able to reach the remote resource by IP address.

For pure NetBIOS name resolution, use the **nbtstat -a ComputerName** command. Net commands, like **net use**, will do a simultaneous NetBIOS and DNS query. Nslookup cannot be used to troubleshoot NetBIOS name resolution, because it will return a host name, not a NetBIOS name.

Resolving a Host or Domain Name to an IP Address

Host names can be directly resolved by the DNS client resolver cache, which contains the entries in the Hosts file, or by a DNS server. Problems here usually involve a misconfigured DNS server, a misspelled Hosts file entry or incorrect IP address, or multiple entries for a single host in the Hosts file. Use Nslookup or Netdiag to diagnose host or domain resolution problems.

Determining the Next-Hop IP Address and Interface

All computers running any version of Windows and the supplied TCP/IP protocol use an IP routing table. The routing table is used to determine the next-hop IP address and interface. The IP routing table stores information about destinations and how they can be reached. There are a series of default entries based on the configuration of the node. You can add entries with TCP/IP utilities, such as the Route command-line tool, or entries can be added dynamically through interaction with routers.

When an IP packet is forwarded, the IP routing table is used to determine the following:

- The next-hop IP address

For a direct delivery (where the destination is a neighboring node), the next-hop IP address is the destination address in the packet. For an indirect delivery (where the destination is not a neighboring node), the next-hop address is the address of a router.

- The next-hop interface

The next-hop interface identifies either a physical interface (for example, a network adapter) or a logical interface (for example, a tunneling interface) that is used to forward the packet.

After the next-hop address and interface are determined, the packet is passed to the ARP. For LAN technologies such as Ethernet and Token Ring, ARP attempts to resolve the MAC address for the next-hop address and forward the packet by using the next-hop interface.

Contents of an IP Routing Table

The following are the fields of a typical IP routing table entry.

Destination The destination can be either an IP address or a class-based subnetted or supernetted network ID. In the Windows Server 2003 IP routing table, this column is named Network Destination.

Network Mask The bit mask that is used to match a destination IP address to the value in the Destination field. In the Windows Server 2003 IP routing table, this column is named Netmask.

Next-Hop The IP address to which the packet is forwarded. In the Windows Server 2003 IP routing table, this column is named Gateway.

Interface The network interface that is used to forward the IP packet.

Metric A number used to indicate the cost of the route so that the best route, among multiple routes to the same destination, can be selected. A common use of the metric is to indicate the number of hops (that is, the number of links or routers to cross) en route to the destination.

Routing Table Entry Types

Routing table entries can be used to store the following types of routes.

Directly attached network routes Routes for subnets to which the node is directly attached. For directly attached network routes, the Next-Hop field can either be blank or contain the IP address

of the interface on that subnet. If the address is local, delivery requires little additional effort. ARP resolves the IP address to a hardware address, typically a MAC address for the destination Ethernet card. The problems found here are typically problems with the ARP cache (such as duplicate addresses) or the subnet mask, and can be solved by using the Arp or Ipconfig tools.

Remote network routes Routes for subnets that are available across routers and are not directly attached to the node. For remote network routes, the Next-Hop field is the IP address of a local router. If the address is remote, the next step is to determine which gateway to use to reach the remote address. In a network with only a single router acting as an external connection, the problem is relatively straightforward. However, in any network with more than one router attached, determining which gateway to use is more difficult.

IP solves the problem by consulting its routing table. This routing table serves as a decision tree that enables IP to decide which interface and which gateway it should use to send the outgoing traffic. The routing table contains many individual routes; each route consists of a destination, net mask, gateway interface, and metric.

The routing table is parsed from the most specific to the most general, so the packet is sent to the first gateway whose routing table entry matches the packet's destination. If two routes are identical, the route with the lowest metric is chosen over the route with a higher metric. In the case of a tie, the node arbitrarily selects which routing table entry to use. Problems found here are addressed with the Route tool or with network configuration changes.

Host routes A route to a specific IP address. Host routes allow routing to occur on a per-IP address basis. For host routes, the network ID is a specific IP address and the network mask is 255.255.255.255.

Default route The route used when a more specific network or host route is not found. The default route destination is 0.0.0.0 with the network mask of 0.0.0.0. The next-hop address of the default route is typically the default gateway of the node.

Route Determination Process

To determine which routing table entry is used for forwarding, IP uses the following process:

- For each entry in the routing table, a bit-wise logical AND operation is performed between the destination IP address and the Network Mask field. The result is compared with the Destination field of the entry for a match.

To perform a bit-wise logical AND between the destination IP address and the network mask of the route, IP compares each bit in the destination IP address to the corresponding bit in the subnet mask. If both bits are 1s, the resulting bit is 1; otherwise, the result is 0. Because of the way in which the subnet mask is defined, the result of the bit-wise logical AND operation is as follows:

- For each bit in the subnet mask that is set to 1, the corresponding bit in the result is copied from the destination IP address.
- For each bit in the subnet mask that is set to 0, the corresponding bit in the result is set to 0.

A good example of performing a bit-wise logical AND is in determining the IP network ID for an IP address configuration. To determine the IP network ID, a bit-wise logical AND of the assigned IP address with its subnet mask is performed. The result is the IP network ID.

For example, for the IP address 192.168.98.112 with the subnet mask 255.255.255.0, the result of the bit-wise logical AND is as follows:

- For the first 24 bits, which correspond to the “255.255.255” portion of the subnet mask, the corresponding bit from the destination IP address is copied, resulting in 192.168.98 for the first three octets.
- For the last 8 bits, which correspond to the “0” portion of the subnet mask, the corresponding bit is set to 0, resulting in 0 for the last octet.

Therefore, 192.168.98.112 AND 255.255.255.0 is 192.168.98.0.

- The list of matching routes is compiled. The route that has the longest match (that is, the route with the highest number of bits set to 1 in the subnet mask) is selected. The longest matching route is the most specific route to the destination IP address. If there are multiple longest match routes (for example, multiple routes to the same network ID), the router uses the lowest metric to select the best route. If there are multiple longest matching routes with the lowest metric, the node arbitrarily selects which routing table entry to use.

The result of the route determination process is the selection of a single route in the routing table. If this process fails to select a route, IP indicates a routing error. For a sending host, an IP routing error is indicated internally to an upper-layer protocol, such as TCP or Usergram Data Protocol (UDP). For a router, an “ICMP Destination Unreachable-Host Unreachable” message is sent to the sending host and the packet is discarded.

Next-Hop Address and Interface Determination Process

After determining the single route in the routing table with which to forward the packet, the next-hop address and interface are determined by the following process:

- If the address in the Next-Hop field is either blank or is an address that is assigned to an interface on the forwarding node, the following occurs:
 - The next-hop address is set to the destination IP address of the IP packet.
 - The next-hop interface is set to the interface that is specified in the Interface field.
- If the address in the Next-Hop field is not an address that is assigned to an interface on the forwarding node, the following occurs:
 - The next-hop address is set to the address in the Next-Hop field for the route.
 - The next-hop interface is set to the interface that is specified in the Interface field.

Example Windows Server 2003 IP Routing Table

The following table lists the default routing table for a Windows Server 2003–based host (that is, not a router) that does not have the IPv6 protocol installed. The host has a single network adapter and is configured with the IP address 157.60.136.41, subnet mask 255.255.252.0 (/22), and a default gateway of 157.60.136.1. To view the IP routing table on a computer running Windows

Server 2003, type **route print** or **netstat -r** at the command prompt. Output similar to the following will be displayed:

```

=====
Interface List
0x1 ..... MS TCP Loopback interface
0x10003 ...00 b0 d0 e9 41 43 ..... 3Com 3C920 EtherLink PCI
=====
Active Routes:
Network Destination        Netmask          Gateway          Interface        Metric
0.0.0.0                    0.0.0.0          157.60.136.1     15.60.136.41     1
127.0.0.0                  255.0.0.0        127.0.0.1        127.0.0.1        1
157.60.136.0              255.255.252.0    157.60.136.41    157.60.136.41    1
157.60.136.41             255.255.255.255  127.0.0.1        127.0.0.1        1
157.60.255.255            255.255.255.255  157.60.136.41    157.60.136.41    1
224.0.0.0                 240.0.0.0        157.60.136.41    157.60.136.41    1
255.255.255.255          255.255.255.255  157.60.136.41    157.60.136.41    1

Default Gateway:          157.60.136.1
=====
Persistent Routes:
None

```

Notice that two interfaces are listed. One interface corresponds to an installed network adapter (3Com EtherLink PCI) and the other is an internal loopback interface (MS TCP Loopback Interface).

The Windows Server 2003 IP routing table uses an IP address to identify an interface in the Interface field for the route. Therefore, the following process determines the next-hop address and interface:

- If the address in the Gateway field is an address that is assigned to an interface on the forwarding node, the following occurs:
 - The next-hop address is set to the destination IP address of the IP packet.
 - The next-hop interface is set to the interface to which the address in the Interface field is assigned.
- If the address in the Gateway field is not an address that is assigned to an interface on the forwarding node, the following occurs:
 - The next-hop address is set to the address in the Gateway field.
 - The next-hop interface is set to the interface to which the address in the Interface field is assigned.

Windows Server 2003 IP routing table entries

The example Windows Server 2003 IP routing table contains the following entries:

- The first entry, network destination of 0.0.0.0 and network mask (netmask) of 0.0.0.0 (/0), is the default route. Any destination IP address that is bit-wise logically ANDed with 0.0.0.0 results in 0.0.0.0. Therefore, the default route is a match for any IP address. If the default route is the longest matching route, the next-hop address is 157.60.136.1 and the next-hop interface is the network adapter that is assigned the IP address 157.60.136.41.
- The second entry, network destination of 127.0.0.0 and netmask of 255.0.0.0 (/8), is the loopback network route. For all packets that are sent to an address of the form 127.x.y.z, the next-hop address is set to 127.0.0.1 (the loopback address) and the next-hop interface is the interface that is assigned the address 127.0.0.1 (the loopback interface).
- The third entry, network destination of 157.60.136.0 and netmask of 255.255.252.0 (/22), is a directly attached network route. If this route is the longest matching route, the next-hop address is set to the destination address in the packet and the next-hop interface is set to the network adapter that is assigned the IP address 157.60.136.41.
- The fourth entry, network destination of 157.60.136.41 and netmask of 255.255.255.255 (/32), is a host route for the IP address of the host. For all IP packets sent to 157.60.136.41, the next-hop address is set to 127.0.0.1 and the next-hop interface is the loopback interface.
- The fifth entry, network destination of 157.60.255.255 and netmask of 255.255.255.255 (/32), is a host route that corresponds to the all-subnets directed broadcast address for the class B network ID 157.60.0.0 (/16). For all IP packets sent to 157.60.255.255, the next-hop address is set to 157.60.255.255 and the next-hop interface is the network adapter that is assigned the IP address 157.60.136.41.
- The sixth entry, network destination of 224.0.0.0 and netmask of 224.0.0.0 (/3), is a route for multicast traffic that is sent by this host. For all multicast packets, the next-hop address is set to the destination address and the next-hop interface is set to the network adapter that is assigned the IP address 157.60.136.41.
- The seventh entry, network destination of 255.255.255.255 and netmask of 255.255.255.255 (/32), is a host route that corresponds to the limited broadcast address. For all IP packets sent to 255.255.255.255, the next-hop address is set to 255.255.255.255 and the next-hop interface is the network adapter that is assigned the IP address 157.60.136.41.

Determining next-hop addresses by using the routing table

The following are examples of how the example routing table would be used to determine the next-hop IP address and interface for several different destinations:

- Unicast destination 157.60.136.48
The longest matching route is the route for the directly attached network (157.60.136.0/22). The next-hop IP address is set to the destination IP address (157.60.136.48), and the next-hop interface is set to the network adapter that is assigned the IP address 157.60.136.41.
- Unicast destination 192.168.0.79
The longest matching route is the default route (0.0.0.0/0). The next-hop IP address is set to the default gateway address (157.60.136.1), and the next-hop interface is the network adapter that is assigned the IP address 157.60.136.41.
- Multicast destination 224.0.0.1

The longest matching route is the 224.0.0.0/3 route. The next-hop IP address is set to the destination IP address (224.0.0.1), and the next-hop interface is the network adapter that is assigned the IP address 157.60.136.41.

- Subnet broadcast destination 157.60.139.255

The longest matching route is the route for the directly-attached network (157.60.136.0/22). The next-hop IP address is set to the destination IP address (157.60.139.255), and the next-hop interface is set to the network adapter that is assigned the IP address 157.60.136.41.

- Unicast destination 157.60.136.41

The longest matching route is the host route for the locally assigned IP address (157.60.136.41/32). The next-hop IP address is set to the destination address (157.60.136.41), and the next-hop interface is set to the loopback adapter.

Troubleshooting Overview

The following sections include information about identifying TCP/IP communication and configuration problems and provide steps that can be taken to correct those problems.

TCP/IP troubleshooting generally follows a set pattern:

1. Verify that the interface on the problem computer is not in a media-disconnected state.
5. Verify that the problem computer's TCP/IP configuration is correct.
6. Verify that a routing path exists between the problem computer and its destination.
7. If link reliability is in question, use Pathping at different times of the day and record the success rate.

If these steps fail to disclose the cause of a problem, use a protocol analyzer, such as Microsoft Network Monitor, to capture the network traffic. For information about Network Monitor, see *Microsoft® Windows® Server 2003 Administrator's Companion* by Microsoft Press.

When troubleshooting a TCP/IP communication problem, ask yourself the following questions:

- Can other computers on the same subnet reach the targeted resource?
- Which applications are failing, which are working, and is there any relationship between them or any pattern you can deduce?
- Is the problem basic IP connectivity, or is it name resolution? If the problem is name resolution, does the failing application use NetBIOS names or host names?
- Did any of the failing applications ever work successfully on this computer in the past?
- Can you identify changes within your computer or network between the time that these applications worked successfully and when they began to fail?

Review the location and timing of the problem to help narrow the scope of the problem. In addition, you can examine TCP/IP failures systematically by referring to the sequence of steps that TCP/IP uses in order to establish communications, as described in "TCP/IP Communication Process Overview" earlier in this chapter.

Unable to Reach an IP Address

Although the TCP/IP stack that is included in Windows Server 2003 is the most reliable version of TCP/IP that has been included in any version of Windows, several things can prevent TCP/IP communications from being successfully established, for example, faulty cabling and hardware, or misconfigured network settings. This section describes several tools available to detect and correct faulty hardware and to disclose errant network configurations.

Checking Network Connections Media-disconnected state

Windows Server 2003 MediaSense provides automatic detection and notification of disconnected or damaged media. Although disconnected media is not strictly a TCP/IP problem, it will stop TCP/IP communications.

If a network cable is unplugged or damaged, it will be detected by MediaSense and the icon for that network connection will be displayed with a red X in the Network Connections folder and in the notification area, at the far right of the taskbar. The red X will also be displayed on network connections if the network hub to which they are connected becomes damaged or unplugged from its power source. As a first step to connectivity troubleshooting, check to see if any network connection icons have a red X, indicating that the connection is in a “media-disconnected state.”

If, for example, your multihomed computer has multiple network adapters connected to one network hub, and all of the network connection icons for those adapters have a red X, the hub to which they are connected might not be functioning properly. If multiple computers are connected to a common hub, and all their network connections are displayed with a red X, this too could be an indication that the hub to which they are connected is not functioning properly. Conversely, if only one of two adapters connected to a common hub is displayed with a red X, it is more likely the cable connecting that adapter to the hub is not plugged in correctly, or the cable has become damaged.

Additionally, check to see whether the connection has been disabled.

Using the Repair Feature

After you have verified that your adapter is not in a media-disconnected or disabled state, you can use the Repair feature to attempt to recover from common network conditions. Repair refreshes network settings. Right-click the icon for a network connection that appears in either the Network Connections folder or in the notification area and then click **Repair** on the shortcut menu. Table 12.1 lists the commands performed by Repair and their command-line equivalents.

Table 12.1 Command Sequence for the Repair Command

Procedure sequence		Command-line equivalent
1	Checks whether DHCP is enabled, and if it is, issues a broadcast renew to refresh the IP address.	Performs similar to ipconfig /renew (See Note below this table)
2	Flushes the ARP cache.	arp -d *
3	Flushes the NETBios cache.	nbtstat -R

4	Flushes the DNS client resolver cache.	ipconfig /flushdns
5	Reregisters with WINS.	nbtstat -RR
6	Reregisters with DNS.	ipconfig /registerdns



Note

A broadcast renew (Repair) causes a computer to accept any lease from any available DHCP server. By contrast, a unicast renew (ipconfig /renew) only renews the existing lease from the last DHCP server the client got a lease from. A broadcast renew is preferable for the cases when a problem arises with the DHCP server from which the client computer last got its lease.

Check Network Configuration

If the computer continues to experience connectivity problems after you have confirmed that your hardware is not in a media-disconnected state and you have run the Repair command, check your network and hardware configuration settings by displaying the Status menu command for a network connection or by using Netdiag.exe or the command-line tool ipconfig. While there is some redundancy in the information reported by these three utilities, each has distinct differences.

Using the Status Command

By using the Status option for a network connection, you can quickly gain access to a variety of configuration settings and statistics for that connection. The information is presented as an extension on the UI of a network connection. To check the status of a connection, open **Network Connections**, right-click the connection, and then click **Status**.

Network Connection Status - General Tab

- The connected or disconnected state of dial-up, wireless, high-speed Internet, or RAS connections.
- The duration of a LAN, dial-up, wireless, high-speed Internet, or RAS connection.
- The speed at which you initially connected.
- For local area connections, the number of bytes transmitted and received during a connection. For other types of connections, the number of bytes transmitted and received during a connection, and the associated compression and error statistics.

Network Connection Status - Support Tab

By switching from the General tab to the Support tab within Status of a connection, you can view the following configuration details:

- Address Type
- IP Address
- Subnet Mask Default Gateway

Network Connection Status - Support Tab Details

Clicking **Details**, on the Support tab, will provide a more granular view of configuration information, consisting of:

- Physical address
- IP address
- Subnet mask
- Default gateway
- DHCP server
- Lease obtained
- Lease expires
- DNS servers
- WINS servers

The advantages of using the **Status** command to check configuration and statistics are:

- You monitor network activity in real time, so you can determine whether or not traffic is actually being processed by the network adapter.
- You can easily access the properties of the connection by clicking **Properties**, and then you can review or change Client, Service, or Protocol network components.

The disadvantages of using this method to check configuration and statistics are:

- You cannot access status for a LAN connection that is in a media-disconnected state.
- The configuration details provided are very limited. For example, no DNS suffixes are reported.
- Status renders results on a per-connection basis only.
- Only IPv4 results are displayed, but not IPv6.
- You can neither print the results from Status, nor redirect the results to be saved as a .txt or .doc file.

Using Ipconfig

When you use the **ipconfig /all** command, it produces a detailed configuration report for all interfaces, including any configured remote access adapters. Ipconfig output can be redirected to be saved as a text (.txt) file, or if Word is installed, as a Word (.doc) file. To do so, type the following at the command prompt:

```
ipconfig > directory\file name.extension
```

with the appropriate file extension. For example, the command

```
ipconfig /all > F:\NetTest\ipcfg.txt
```

saves the output as a text file named ipcfg.txt in the directory F:\NetTest\.

The output of Ipconfig can be reviewed to find any problems in the computer network configuration. For example, if a computer has been configured with an IP address that is a duplicate of an existing IP address that has already been detected, the subnet mask appears as 0.0.0.0.

The following example illustrates the results of an **ipconfig /all** command on a multihomed computer that is configured to use a DHCP server for automatic TCP/IP configuration and use WINS and DNS servers for name resolution, and does not have the IPv6 protocol installed:

Windows IP Configuration

```
Host Name . . . . . : testpcl
Primary Dns Suffix . . . . . : contoso.example.com
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : contoso.example.com
                                contoso.com
                                example.com
```

Ethernet adapter Local Area Connection 1:

```
Connection-specific DNS Suffix . . : corp.example.com
Description . . . . . : PCI 100 Ethernet Adapter
Physical Address. . . . . : 00-02-B3-22-01-5D
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . . : Yes
IP Address. . . . . : 172.16.48.10
Subnet Mask . . . . . : 255.255.248.0
Default Gateway . . . . . : 172.16.48.03
DHCP Server . . . . . : 157.54.8.118
DNS Servers . . . . . : 172.16.14.119
                        172.56.236.138
Primary WINS Server . . . . . : 172.16.48.04
Secondary WINS Server . . . . . : 172.16.48.05
Lease Obtained. . . . . : Thursday, July 10, 2003 1:49:40 PM
Lease Expires . . . . . : Friday, July 18, 2003 1:49:40 PM
```

Ethernet adapter Local Area Connection 2:

```
Connection-specific DNS Suffix . . :
Description . . . . . : PCI Fast Ethernet Adapter
(Generic)
Physical Address. . . . . : 00-00-F8-03-6F-D3
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . . : Yes
Autoconfiguration IP Address. . . . : 169.254.225.167
Subnet Mask . . . . . : 255.255.0.0
Default Gateway . . . . . :
```

The advantages of using Ipconfig are:

- Information can be provided details for multihomed computers in a single operation.
- The results from Ipconfig can be redirected and saved as a .txt file, or if Word is installed on the computer, as a.doc file.
- Ipconfig displays results for both IPv4 and IPv6.
- Ipconfig provides more information than what is provided by the **Status** command for a connection.

The disadvantage of using Ipconfig to check configuration is:

- You cannot use Ipconfig to display the configuration of a remote computer.

Using Netdiag.exe

The Netdiag.exe tool isolates networking and connectivity problems, performing a more extensive series of tests than Ipconfig. These tests and the key network status information they expose can help you identify and isolate network problems. Because this tool does not require that parameters or options be specified, you can focus on analyzing the output, rather than training users on tool usage.

You can install Netdiag.exe and additional support tools by running Suptools.msi, located in the Support\Tools\ folder on the Windows Server 2003 product CD. Double-click **Suptools.msi** to start the Windows Support Tools Setup Wizard.



Note

By default, the **Destination Directory** page of the wizard will install the support tools in *systemroot*\Program Files\Support Tools\. Because Netdiag.exe must be run from the command line, you might want to designate a different installation location in order to shorten the path that you enter every time you run the tool. For example, during installation you can use the wizard's **Browse** button to designate the folder that you created, D:\Tools.

To install Netdiag.exe independently from other Support Tools, double-click the **Support\Tools\Support.cab** file on your Windows Server 2003 product CD. Double-click the **Netdiag.exe** file to start the extraction, and then follow the provided instructions.

To run Netdiag.exe from the command line, provide the path to the location of the extracted Netdiag.exe file. For example, if you extract Netdiag.exe to a folder you created and named E:\Tools, type the following at the command prompt:

```
e:\tools\netdiag.exe
```

Or you can change to the directory E:\Tools and then type the command **Netdiag.exe**.

The output of Netdiag.exe can be redirected to a folder and saved as a .txt or .doc file. To do so, type the following at the command prompt:

```
netdiag.exe > directory\file name
```

with the appropriate file name and extension. For example, type the following at the command prompt:

```
E:\Tools\Netdiag.exe > E:\NetTest\Ntdiag.doc
```

Netdiag will be run from the directory E:\Tools and the results will be saved in a file named E:\NetTest\ as the Word document Ntdiag.doc.

The output of Netdiag.exe will provide a detailed list of configuration information and the tests performed. In addition to confirming correct configuration of IP address, subnet mask, default gateway, DNS servers, and WINS servers, you should review the output of Netdiag.exe for reported errors. The results will disclose details about configuration errors that were detected by Netdiag.exe. Investigate entries with “Skipped,” “Failed,” and “WARNING” messages associated with the various tests.

In the following example, Netdiag is run on a computer that is configured to use a DHCP server for automatic TCP/IP configuration, and WINS and DNS servers for name resolution. Output similar to the following is displayed:

```
Computer Name: testpcl
DNS Host Name: testpcl.contoso.example.com
System info : Windows Server 2003
Processor : x86 Family
List of installed hotfixes :
    Q147222

Netcard queries test . . . . . : Passed
    [WARNING] The net card 'PCI Fast Ethernet Adapter (Generic)' may not be
working because it has not received any packets.

Adapter : Local Area Connection 1

Netcard queries test . . . : Passed
Host Name. . . . . : testpcl.example.com
IP Address . . . . . : 172.16.48.10
Subnet Mask. . . . . : 255.255.248.0
Default Gateway. . . . . : 172.16.48.03
Primary WINS Server. . . . : 172.16.48.04
Secondary WINS Server. . . : 172.16.48.05
Dns Servers. . . . . : 172.16.14.119
                    172.56.236.135

AutoConfiguration results. . . . . : Passed

Default gateway test . . . : Passed

NetBT name test. . . . . : Passed
    [WARNING] At least one of the <00> 'WorkStation Service', <03> 'Messenger
Service', <20> 'WINS' names is missing.

WINS service test. . . . . : Passed
```

Adapter : Local Area Connection 2

Netcard queries test : Passed

Host Name : testpcl
 Autoconfiguration IP Address : 169.254.225.167
 Subnet Mask : 255.255.0.0
 Default Gateway :
 Dns Servers :

AutoConfiguration results : Failed
 [WARNING] AutoConfiguration is in use. DHCP not available.

Default gateway test : Skipped
 [WARNING] No gateways defined for this adapter.

NetBT name test : Passed
 [WARNING] At least one of the <00> 'WorkStation Service', <03> 'Messenger Service', <20> 'WINS' names is missing.
 No remote names have been found.

WINS service test : Skipped
 There are no WINS servers configured for this interface.

Global results:

Domain membership test : Passed

NetBT transports test : Passed
 List of NetBt transports currently configured:
 NetBT_Tcpip_{0B19AD54-2CA7-4795-8729-FE7494F2316A}
 NetBT_Tcpip_{C3C96E7E-4C54-4C87-9462-67D21D3E3D74}
 2 NetBt transports currently configured.

Autonet address test : Passed

IP loopback ping test : Passed

Default gateway test : Passed

NetBT name test : Passed
 [WARNING] You don't have a single interface with the <00> 'WorkStation Service', <03> 'Messenger Service', <20> 'WINS' names defined.

Winsock test : Passed

DNS test : Passed
 [WARNING] Cannot find a primary authoritative DNS server for the name

```
'testpcl.contoso.example.com.'. [ERROR_TIMEOUT]
The name 'testpcl.contoso.example.com.' may not be registered in DNS.

Redir and Browser test . . . . . : Passed
List of NetBt transports currently bound to the Redir
NetBT_Tcpip_{0B19AD54-2CA7-4795-8729-FE7494F2316A}
NetBT_Tcpip_{C3C96E7E-4C54-4C87-9462-67D21D3E3D74}
The redir is bound to 2 NetBt transports.

List of NetBt transports currently bound to the browser
NetBT_Tcpip_{0B19AD54-2CA7-4795-8729-FE7494F2316A}
NetBT_Tcpip_{C3C96E7E-4C54-4C87-9462-67D21D3E3D74}
The browser is bound to 2 NetBt transports.

DC discovery test. . . . . : Passed

DC list test . . . . . : Passed

Trust relationship test. . . . . : Passed
Secure channel for domain 'contoso' is to '\\contoso-dc-
02.contoso.example.com'.

Kerberos test. . . . . : Passed

LDAP test. . . . . : Passed

Bindings test. . . . . : Passed

WAN configuration test . . . . . : Skipped
No active remote access connections.

Modem diagnostics test . . . . . : Passed

IP Security test . . . . . : Skipped
```

Note: run "netsh ipsec dynamic show /?" for more detailed information

The command completed successfully

The advantages of using Netdiag.exe are:

- Information is provided for multihomed computers in a single operation, as compared to the Network Connection Status, which is provided on a per-connection basis.
- The results from Netdiag.exe can be redirected and saved as a .txt or .doc file. Network Connection Status provides no method to redirect the output to a .txt or .doc file.
- Netdiag.exe will display results for both IPv4 and IPv6.
- Netdiag.exe runs a wider variety of tests resulting in a more detailed diagnostic report than either Ipconfig or Network Connection Status.

Test Network Connection by Using Ping and Pathping

If no problems appear in the TCP/IP configuration, the next step is to use the Ping and Pathping tools to test the computer's ability to connect to other host computers on the TCP/IP network.

Ping is a tool that helps to verify IP-level connectivity; Pathping is a tool that detects packet loss over multiple-hop trips. The **ping** command sends a series of Internet Control Message Protocol (ICMP) Echo messages to a destination based on the destination's host name or IP address. Use Ping whenever you want to verify that a host computer can send IP packets to a destination host. You can also use the Ping tool to isolate network hardware problems and incompatible configurations.



Note

If you run **ipconfig /all** and the IP address configuration is properly displayed, there is no need to ping the loopback address or the local computer's IP address — **ipconfig** has already done the equivalent in order to display the host's configuration.

It is best to verify that a route exists between the local computer and a network host by first using Ping and the IP address of the network host to which you want to connect. The command syntax is:

```
ping IPAddress
```

Perform the following steps when using Ping:

1. Ping the loopback address to verify that TCP/IP is installed and configured correctly on the local computer.

```
ping 127.0.0.1
```

8. Ping the IP address of the local computer to verify that it was added to the network correctly. Note that if the routing table is correct, this simply forwards the packet to the loopback address of 127.0.0.1.

```
ping IPAddressOfLocalHost
```

9. Ping the IP address of the default gateway to verify that the default gateway is functioning and that you can communicate with a local host on the local network.

```
ping IPAddressOfDefaultGateway
```

10. Ping the IP address of a remote host to verify that you can communicate through a router.

```
ping IPAddressOfRemoteHost
```

11. Ping the host name of a remote host to verify that you can resolve a remote host name.

```
ping HostNameOfRemoteHost
```

12. Run a Pathping analysis to a remote host to verify that the routers in the path to the destination are operating correctly.

```
pathping IPAddressOfRemoteHost
```

**Note**

If your local address is returned as 169.254.y.z, you have been assigned an IP address by the Automatic Private IP Addressing (APIPA) feature of Windows Server 2003. This means that the local DHCP server is not configured properly or cannot be reached from your computer, and an IP address has been assigned automatically with a subnet mask of 255.255.0.0. Enable or correct the DHCP server, restart the local computer, and see if the networking problem persists.

If your local address is returned as 0.0.0.0, MediaSense detected that the network adapter is not connected to a network. To correct this problem, ensure that the network adapter and network cable are connected to a working hub. If the connection is solid, reinstall the network adapter's drivers or install a new network adapter.

Ping uses host name resolution to resolve a computer name to an IP address, so if pinging by address succeeds, but pinging by name fails, the problem lies in host name resolution, not network connectivity. For more information about troubleshooting host name resolution, see the section "Unable to Reach a Host or NetBIOS Name" in this chapter.

If you cannot use Ping successfully at any point, check the following:

- The local computer's IP address is valid and is displayed correctly on the **General** tab of the **Internet Protocol (TCP/IP) Properties** dialog box and on the **Support** tab of the **Status** command for the network connection, or when using either the Ipconfig or Netdiag.exe tools.
- A default gateway is configured and the link between the host and the default gateway is operational. You should get a reply when you ping the address of the default gateway. Note that Windows Server 2003 uses Dead Gateway Detection. If your computer is configured with multiple gateways and a connection routed through the default gateway attempts to send a TCP packet to the destination a number of times without receiving a response, the next-hop address for the TCP connection is changed to use the next default gateway in the list.

**Important**

If the remote system being pinged is across a high-delay link, such as a satellite link, responses might take longer to be returned. The **-w** (wait) option can be used to specify a longer time-out. The following example shows a **-w** option set to wait 2 seconds (2,000 milliseconds) for a response before timing out, a **-n** (number of pings) option set to 2 Echo messages, and a **-l** option specifying each ping to be 1450 bytes in size. The example pings the address 172.16.48.10.

```
C:\>ping -w 2000 -n 2 -l 1450 172.16.48.10
```

The following output is displayed:

```
Pinging 172.16.48.10 with 1450 bytes of data:
```

```
Reply from 172.16.48.10: bytes=1450 time=1542ms TTL=32
```

```
Reply from 172.16.48.10: bytes=1450 time=1787ms TTL=32
```

```
Ping statistics for 172.16.48.10:  
Packets: Sent = 2, Received = 2, Lost = 0 (0% loss),  
Approximate round trip times in milli-seconds:  
Minimum = 0ms, Maximum = 1787ms, Average = 1664ms
```

Flush the ARP Cache

Incorrect entries in your ARP cache can prevent connectivity to local hosts or remote hosts (if the ARP cache entry for the default gateway is incorrect). To display the contents of the ARP cache, use the **arp -a** (or **arp -g**) command. To flush the ARP cache, use **arp -d ***. The results from **arp -a** (or **arp -g**) can be redirected to save as a .txt file, or, if Word is installed, as a Word (.doc) file, by using the syntax **arp -a > Directory\FileName**, with the appropriate .txt or .doc file name extension. Before you change the configuration of the ARP cache, it is a good idea to record configuration settings by redirecting the contents of the ARP cache to a .txt or .doc file.

Verify Default Gateway

Check the default gateway. The gateway address must be on the same subnet as the local host; if it is not, packets from the host computer cannot be forwarded to any location outside the local subnet. Then check to make sure that the default gateway address is correctly configured on the host, either through automatic or manual configuration.

Ping Remote Host

If the default gateway responds correctly, ping a remote host to ensure that remote network communications are operating as expected. If this fails, use the Tracert tool to examine the path to the destination. For IP routers that are computers running Microsoft® Windows NT, Windows 2000, or Windows Server 2003, use the **route print** command or Routing and Remote Access to examine the IP routing table. For IP routers that are not running Windows NT operating systems, Windows 2000, or Windows Server 2003, use the appropriate utility or facility to examine the IP routing table.

Ping commonly returns four error messages during troubleshooting.

TTL Expired in Transit

The number of hops required to reach the destination exceeds the TTL (Time to Live) set by the sending host to forward the packets. The default TTL value for ICMP Echo messages sent by Ping is 128. If this is not enough to travel the required number of links to a destination, you can increase the TTL by using **ping -i**, up to a maximum of 255 links. If increasing the TTL value fails to resolve the problem, the packets are being forwarded in a routing loop — that is, a circular path among routers. Use Tracert to track down the set of routers in the routing loop, which appears as a repeated series of the same IP addresses in the Tracert report. Next, make an appropriate change to the routing tables of the routers in the routing loop, or inform the administrator of the remote router of the problem.

Destination Host Unreachable

This message indicates one of two problems: either the local system has no route to the desired destination, or a remote router reports that it has no route to the destination. The form of the message can distinguish the two problems:

- If the message is simply “Destination Host Unreachable,” there is no route from the local system and the packets to be sent were never put on the wire. Use the Route utility to check the local routing table to see if the route to the destination is errant or missing.
- If the message is “Reply From *IP Address*: Destination Host Unreachable,” the routing problem occurred at a remote router, whose address is indicated by *IP Address*. Use the appropriate tool to check the IP routing table of the router assigned the IP address of *IP Address*.

If you pinged by using an IP address, retry it with a host name to ensure that the IP address you tried is correct.

Request Timed Out

This message indicates that no Echo Reply messages were received within the default time of four seconds. This can be due to many different causes; the most common include network congestion, failure of ARP to resolve the next-hop MAC address, packet filtering, routing error, or a silent discard. Most often, it means that a route back to the sending host has failed. This might be because the destination host does not know the route back to the sending host, or one of the intermediary routers does not know the route back, or even that the destination host’s default gateway does not know the route back. Check the routing table of the destination host to see whether it has a route to the sending host before you check the routing tables at the routers.

If the remote routing tables are correct and contain a valid route back to the sending host, see if the ARP cache lacks the proper address by using the **arp -a** command to print the contents of the ARP cache. Also, check the subnet mask to be sure that a remote address has not been interpreted as local.

Next, use Tracert to determine the path to the destination. Although Tracert does not record the path that the Echo Reply messages follow on their return path, it might show that the packet made it to the destination. If this is the case, the problem is probably a routing issue on the return path. If the trace does not quite reach the destination, it might be because the target host is protected by a firewall. When a firewall protects the destination, ICMP packet filtering prevents the ping packets — or any other ICMP messages — from crossing the firewall and reaching their destination.

To check for network congestion, simply increase the allowed latency by setting a higher wait time, such as 5,000 milliseconds, by using the **ping -w** command. Try to ping the destination again. If the request still times out, congestion is not the problem.

Unknown Host

This error message indicates that the requested host name cannot be resolved to its IP address; check that the name is entered correctly and that the DNS servers can resolve it.

Test IP-to-MAC Address Resolution with ARP

TCP/IP for Windows Server 2003 allows applications to communicate over a network with another computer by using an IP address, a host name, or a NetBIOS name. However, regardless of which naming convention is used, the next-hop address for the destination must ultimately be resolved to a hardware address — also known as a media access control or MAC address — for shared access media such as Ethernet and Token Ring.

ARP allows a host to find the MAC address of the next-hop IP address on the same physical network. To make ARP efficient, each computer caches IP-to-MAC address mappings to eliminate repetitive ARP broadcast requests.

The Arp tool allows a user to view and modify ARP table entries on the local computer. The **arp** command is useful for viewing the ARP cache and resolving address resolution problems.

A static entry can be added to an ARP file by using the **arp -s IPAddress MACAddress** command. However, be careful when you add such static ARP cache entries, because it is easy to enter the wrong MAC address for an IP address. Note that static ARP entries are cleared when the computer is restarted.

Detecting Duplicate IP Addresses by Using ARP

When starting up, Windows performs a gratuitous ARP to detect any duplication with its own IP address. A gratuitous ARP is an ARP request for a node's own IP address. If a node sends an ARP Request frame for its own IP address and no ARP Reply frames are received, the node determines that no other nodes are using its assigned IP address. Although this detects most cases of duplicate IP addresses, in a few situations two TCP/IP hosts (either Microsoft or non-Microsoft) on the same network can be configured for the same IP address.

The MAC and IP address mapping is done by the ARP module, which uses the first ARP response it receives. Therefore, the reply from the impostor computer will sometimes come back before the reply from the intended computer.

Use the **arp -a** command to display the mappings in the ARP cache. If you know the MAC address for the remote computer you want to use, you can easily determine whether the two match. If they do not, use the **arp -d** command to delete the entry, then use Ping with the same address (forcing an ARP), and check the MAC address in the cache again by using **arp -a**.

If both computers are on the same network, you will eventually get a response from the impostor computer. If not, you might have to capture the traffic from the impostor host with Network Monitor to determine the owner or location of the system. For more information about Network Monitor, see *Microsoft® Windows® Server 2003 Administrator's Companion*.

Detecting Invalid Entries in the ARP Cache

Troubleshooting the ARP cache can be one of the more difficult tasks in network administration because the problems associated with it are frequently intermittent.

The exception to this rule is when you find that the wrong host responds to an ARP request, creating an invalid entry in the ARP cache. The symptoms of invalid entries in the ARP cache are harder to reproduce and involve intermittent problems that only affect a few hosts. The underlying problem is that two computers are using the same IP address on the network. You only see the problems intermittently because the most recent ARP table entry is always the one from the host that responded more quickly to any particular ARP request.

To address the problem, display the ARP table. Type the following at the command prompt:

```
C:\>arp -a 172.16.0.142
```

Output similar to the following is displayed:

```
Interface: 172.16.0.142
Internet Address      Physical Address      Type
172.16.0.1           00-e0-34-c0-a1-40     dynamic
172.16.1.231         00-00-f8-03-6d-65     dynamic
172.16.3.34          08-00-09-dc-82-4a     dynamic
172.16.4.53          00-c0-4f-79-49-2b     dynamic
157.59.5.102         00-00-f8-03-6c-30     dynamic
```

Because addresses assigned by DHCP typically do not cause address conflicts like those described here, the main source of these conflicts is likely to be static IP addresses. Maintaining a list of static addresses (and corresponding MAC addresses) as they are assigned can help you track down any address conflict just by examining the IP and MAC address pairs from the ARP table and comparing them to the recorded values.

If you do not have a record of all IP and MAC address pairs on your network, you might want to examine the manufacturer bytes of the MAC addresses for inconsistencies. The first three bytes of each MAC address identify the card's manufacturer. These three-byte numbers are called Organizationally Unique Identifiers (OUIs) and are assigned by the Institute of Electrical and Electronics Engineers (IEEE). Knowing what equipment you installed and comparing that with the values returned by **arp -a** might allow you to determine which static address was entered in error.

Verify Persistent Routing Table Entries

The next area to examine is the persistent entries in your routing tables. You can view these by using the Route utility. Persistent entries are added by using the **route add -p** command or through the Routing and Remote Access service. Change incorrect entries by using the **route change** command. You can use Routing and Remote Access to add a static route.

► **To add a static route by using Routing and Remote Access**

1. Open Routing and Remote Access.
2. In the console tree, right-click **Static Routes**.

Where?

Routing and Remote Access

ServerName

IP Routing

StaticRoutes

3. Click **New Static Route**.
4. In the **Static Route** dialog box, enter the interface, destination, network mask, gateway, and metric. If this is a demand-dial interface, Gateway is unavailable.

Use Tracert and Pathping

If the routing table configuration is correct, the problem might be with a router or link at any point along the route. You can trace the path to the destination computer by using Tracert and Pathping to pinpoint the problem.

Unless there is only one path to the destination host, be certain to use these tools to map out the route a few times, particularly if you are noticing intermittent packet loss. The datagram might be sent along different paths, and a faulty router might be the problem.

Use Tracert when you have no connectivity to a site under investigation, because it tells you where connectivity stops. Pathping is more useful when you have connectivity to a site but are experiencing some packet loss or high delay. In these cases, Pathping tells you exactly where packet loss is occurring.

Verify Server Services on the Remote Computer

Sometimes a system configured as a remote gateway or router is not functioning as a router. To confirm that the remote computer you want to contact is set up to forward packets, you can either examine it by using a remote administration tool (assuming that it is a computer you administer) or you can attempt to contact the person who maintains the computer.

You can contact the administrator responsible for a remote network by using the databases maintained by InterNIC. The easiest way to do this is to use the Whois tool to find the appropriate person's name and contact information from the InterNIC database. To find the Whois tool, see the [InterNIC Web site](http://go.microsoft.com/fwlink/?linkid=8177) at <http://go.microsoft.com/fwlink/?linkid=8177>.

Check IPSec on the Initiating Host

IP security (IPSec) can increase the defenses of a network, but it can also make changing network configurations or troubleshooting problems more difficult. In some cases, IPSec running on the initiating host of a computer under investigation can create difficulties in connecting to another host. To determine if this is a source of problems, temporarily turn off IPSec by using the **net stop policyagent** command and attempt to run the requested network service or function

If the problem disappears when IPSec policy settings are turned off, you know that the additional IPSec processing burden, or its packet filtering, are responsible for the problem. To solve the

problem, restart the IPSec service by using the **net start policyagent** command, and then use the following procedure.

To see the active filters, type the following command at a command prompt:

```
netdiag /test:ipsec /debug
```

You can optionally redirect the output of this command to a text file so you can view it with a text editor (such as Notepad) by typing the following command:

```
netdiag /test:ipsec /debug > filename.txt
```

► **To assign or unassign IPSec policy for Active Directory–based Group Policy**

1. Open Active Directory Users and Computers. (To open Active Directory Users and Computers, click **Start**, click **Control Panel**, double-click **Administrative Tools**, and then double-click **Active Directory Users and Computers**.)

In the console tree, right-click the domain or organizational unit for which you want to set Group Policy.

Where?

Active Directory Users and Computers [DomainControllerName.DomainName]

Domain

OrganizationalUnit

ChildOrganizationalUnit

2. Click **Properties**, and then click the **Group Policy** tab.
3. Click **Edit** to open the Group Policy object that you want to edit. Or click **New** to create a new Group Policy object, and then click **Edit**.
4. In the Group Policy console tree, click **IP Security Policies on Active Directory**.

Where?

PolicyName [ComputerName] Policy

Computer Configuration

Windows Settings

Security Settings

IP Security Policies on Active Directory

5. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:
 - To assign the policy, on the **Action** menu, click **Assign**.
 - To unassign the policy, on the **Action** menu, click **Unassign**.

► **To assign or unassign IPSec policy for local computer policy**

1. Click **Start**, click **Run**, type MMC, and then click **OK**.
2. Click **File**, click **Add/Remove Snap-in**, and then click **Add**.

3. Click **Group Policy Object Editor**, and then click **Add**.
4. Click **Finish**, click **Close**, and then click **OK**.
5. In the Group Policy console tree, click **IP Security Policies on Local Computer**.

Where?

Local Computer Policy

Computer Configuration

Windows Settings

Security Settings

IP Security Policies on Local Computer

6. In the details pane, click the IPSec policy that you want to assign or unassign, and then do one of the following:
 - To assign the policy, on the **Action** menu, click **Assign**.
 - To unassign the policy, on the **Action** menu, click **Un-assign**.



Note

For more information about IPSec for Windows Server 2003, on your computer running Windows Server 2003, click **Start**, click **Run**, type **hh IPSECconcepts.chm**, and then click **OK**.



Important

An IPSec policy might remain active even after the IPSec policy or Group Policy object to which it is assigned has been deleted. Therefore, you should unassign the IPSec policy before you delete either the policy or the Group Policy object. To prevent problems, unassign the IPSec policy in the Group Policy object. Wait 24 hours to ensure that the change has been propagated, and then delete the IPSec policy setting or Group Policy object.

Check Packet Filtering

Any mistakes in packet filtering at the TCP/IP, router, proxy server, Routing and Remote Access service, or IPSec level can cause address resolution or connectivity to fail. To determine whether TCP/IP filtering is the source of a network problem, you must disable the TCP/IP packet filtering.



To disable TCP/IP packet filtering

1. In Control Panel, double-click **Network Connections**.
2. Right-click the connection, and then click **Properties**.
3. Select **Internet Protocol (TCP/IP)**, and then click the **Properties** tab.

4. Click **Advanced**, and then click the **Options** tab.
5. Under **Optional Settings**, click **TCP/IP Filtering**, and then click **Properties**.
6. Clear the **Enable TCP/IP Filtering (All adapters)** check box, and then click **OK**.

Try pinging an address by using its DNS name, its NetBIOS name, or its IP address. If the attempt succeeds, the packet filtering options might be incorrectly configured or might be too restrictive. For instance, the filtering might permit the computer to act as a Web server, but, in the process, it might disable tools like Ping or remote administration. Restore a wider range of permissible filtering options by changing the permitted TCP, UDP, and IP port values.

If the attempt still fails, another form of packet filtering might be interfering with your networking. For more information about Routing and Remote Access IP packet filtering functions, on your computer running Windows Server 2003, click **Start**, click **Run**, type **hh RRASconcepts.chm**, and then click **OK**. For more information about IPSec packet filtering, on your computer running Windows Server 2003, click **Start**, click **Run**, type **hh IPSECconcepts.chm**, and then click **OK**.

Unable to Reach a Host or NetBIOS Name

TCP/IP for Windows Server 2003 allows applications to communicate over a network with another computer by using three types of destination designations:

- IP address
- Host name
- NetBIOS name

This section describes how to troubleshoot either host name or NetBIOS name resolution problems. Problems with IP addressing are covered in “Unable to Reach an IP Address” earlier in this chapter.

The first step is to determine which applications are failing. Typically, this is Internet Explorer, **net use**, Telnet, or FTP. Making this determination helps with the next step, which is to determine whether the failure is due to a host name problem or NetBIOS name resolution problem.

The easiest way to distinguish host name problems from NetBIOS name resolution problems is to find out whether the failing application uses NetBIOS or Sockets. If it uses Sockets, the problem lies with host name resolution. NetBIOS applications are among the most common, including the various **net** commands, Windows Explorer, and My Network Places. Sockets applications include Internet Explorer and other Web browsers, Telnet, and FTP.

The following sections describe the processes that occur when a host name or a NetBIOS name is used to connect with hosts on a TCP/IP network.

Error 53

The most common symptom of a problem in NetBIOS name resolution is when the **ping** command returns an Error 53 message. The Error 53 message is generally returned when name resolution fails for a particular computer name. Error 53 can also occur when there is a problem establishing a NetBIOS session. To distinguish between these two cases, use the following procedure:

▶ **To determine the cause of an Error 53 message**

At the command prompt, type

```
net view \\HostName
```

where *HostName* is a network resource you know is active.

If this works, your name resolution is probably not the source of the problem. To confirm this, ping the host name, because name resolution can sometimes function properly even when **net use** returns an Error 53 (such as when a DNS or WINS server has a bad entry). If Ping also shows that name resolution fails (by returning the “Unknown host” message), check the status of your NetBIOS session.

▶ **To check the status of your NetBIOS session**

At the command prompt, type

```
net view \\IPAddress
```

where *IP Address* is the same network resource you used in the above procedure.

If this also fails, the problem is in establishing a session.

If the computer is on the local subnet, confirm that the name is spelled correctly and that the target computer is running TCP/IP as well. If the computer is not on the local subnet, be sure that its name and IP address mapping are available in the DNS database, the Hosts or Lmhosts file, or the WINS database.

If all TCP/IP elements appear to be installed properly, use Ping with the remote computer to be sure its TCP/IP protocol is working.

Cannot Connect to Remote Systems by Using Host Name

If the problem is not NetBIOS but Sockets, the problem is related to either a Hosts file entry or a DNS configuration error. To determine why only IP addresses but not host names work for connections to remote computers, make sure that the appropriate Hosts file and DNS setup have been configured for the computer.

Check the Hosts File

The entries in the Hosts file are used to create entries in the DNS client resolver cache, which you can view by using the **ipconfig /displaydns** command. Check the contents of the DNS client resolver cache and the Hosts file for the correct entries.

Check Your DNS Configuration

If you are using DNS, be sure that the IP addresses of the DNS servers are entered correctly in TCP/IP properties, and in the proper order, by using the following procedure.

► **To check DNS configuration**

1. In Control Panel, double-click **Network Connections**.
13. Right-click the connection you want to examine, and then click **Properties**.
14. Click **Internet Protocol (TCP/IP)**, and then click **Properties**.
15. Click **Advanced** in the **Microsoft TCP/IP Properties** dialog box.
16. Click the **DNS** tab.
17. Confirm that DNS is configured properly. If the DNS server IP address is missing, add it to the list of DNS server addresses.

This procedure is for computers on which **Use the following DNS server address** is selected on the General tab of TCP/IP Properties, or statically configured computers; DHCP clients do not have DNS server in the list.

Ping with the remote computer's host name and then with its IP address to determine whether the host address is being resolved properly. If the host name ping fails and the IP address ping succeeds, the problem is with name resolution. You can test whether the DNS servers are running by pinging their IP addresses or by opening a Telnet session to port 53 on the DNS server. If the connection is established successfully, the DNS service is working on the DNS server. For more information about using Telnet, type **Telnet /?** at the command prompt. After you have verified that the DNS service is running, you can use the Nslookup tool to send queries to the DNS server to further verify the status of the records you are looking for.

If pings by both IP address and by name fail, the problem is with network connectivity, such as basic connectivity or routing. For more information about troubleshooting network connectivity, see "Troubleshooting IP Routing" later in this chapter.

Host Name Resolution by Using a DNS Server

DNS is a distributed database that maps domain names to data. A user can query DNS by using hierarchical, friendly names to locate computers and other resources on an IP network. This allows it to largely replace the function that was once performed by the Hosts file. To do so, it resolves friendly names to IP addresses as described below. Sometimes an answer might be provided by any server along the line, preventing the need for further iterative queries.

1. The client contacts the DNS name server with a recursive query for *name.contoso.com*. The server must now return the answer or an error message.

2. The DNS name server checks its cache and zone files for the answer, but does not find it. It contacts a server at the root of the Internet (a root DNS server) with an iterative query for *name.contoso.com*.
3. The root server does not know the answer, so it responds with a referral to an authoritative server in the .com domain.
4. The DNS name server contacts a server in the .com domain with an iterative query for *name.contoso.com*.
5. The server in the .com domain does not know the exact answer, so it responds with a referral to an authoritative server in the contoso.com domain.
6. The DNS name server contacts the server in the contoso.com domain with an iterative query for *name.contoso.com*.
7. The server in the contoso.com domain does know the answer. It responds with the correct IP address.
8. The DNS name server responds to the client query with the IP address for *name.contoso.com*.

This example is specific to the Internet. For more information about DNS host name resolution, recursive queries, and iterative queries, on your computer running Windows Server 2003, click **Start**, click **Run**, type **hh DNSconcepts.chm**, and then click **OK**.

DNS Error Messages

Errors in name resolution can occur if the entries in a DNS server or client are not configured correctly, if the DNS server is not running, or if there is a problem with network connectivity. To determine the cause of any name resolution problem, you can use the Nslookup utility.

Failed queries will return a variety of messages, depending on the nature of the failure. For example, if the following command is used:

```
C:\nslookup DestinationHost
```

and the server cannot resolve the name, the following output is displayed:

```
Server: FullyQualifiedDomainName
Address: ServerIPAddress
*** FullyQualifiedDomainName can't find DestinationHost: Non-existent domain
```

In other cases, the request to the DNS service time out without a reply and returns a message in the following format:

```
C:\nslookup ValidHost
Server: [IPAddress]
Address: w.x.y.z
DNS request timed out.
timeout was 2 seconds.
```

If the server fails to answer the request, Nslookup returns an error message in the following format:

```
C:\nslookup
*** Can't find server name for address IPAddress: No response from server
*** Default servers are not available.
```

This message indicates that the DNS server cannot be reached; it does not indicate why the server is unavailable. The server might be offline, the host computer might not have the DNS service enabled, or there might be a hardware or routing problem.

For more information about DNS troubleshooting, on your computer running Windows Server 2003, click **Start**, click **Run**, type **hh DNSconcepts.chm**, and then click **OK**.

Host Name Resolution by Using a Hosts File

A computer using its Hosts file for name resolution performs the following steps.

1. Computer A enters a command by using the host name of Computer B.
2. Computer A checks the contents of the DNS client resolver cache, which contains the entries in the Hosts file. When the host name of Computer B is found, it is resolved to an IP address. Then the packet is sent by using the normal route determination and address resolution processes.

The following Hosts file problems can cause networking errors:

- The Hosts file does not contain the particular host name.
- The host name in the Hosts file or in the command is misspelled.
- The IP address for the host name in the Hosts file is invalid or incorrect.
- The Hosts file contains multiple entries for the same host on separate lines. Because the Hosts file is parsed from the top, the first entry is used to populate the DNS client resolver cache.

This file is not updated dynamically; all entries are added manually, and the IP address and friendly host name are always separated by one or more spaces or tab characters. The file format is the following:

```
IP Address      Friendly Name
172.16.48.10    testpc1 # Remarks are denoted with a #.
```

If you are having trouble connecting to a remote system by using a host name and are using a Hosts file for name resolution, the problem might be with the contents of that file. Make sure the name of the remote computer is spelled correctly in the Hosts file and by the application that uses it. You can find the Hosts file in *systemroot*\System32\Drivers\Etc.

Check the Lmhosts File

The name resolution problem might be in your Lmhosts file, which looks for addresses for NetBIOS names sequentially from the top down. If more than one address is listed for the same host name, TCP/IP returns the first value it encounters, whether that value is accurate or not.

You can find the Lmhosts file in *systemroot\System32\Drivers\Etc*. This file does not exist by default; a sample file named *Lmhosts.sam* exists. This file must be renamed or copied to *Lmhosts* before it is used.

Although *systemroot\System32\Drivers\Etc* is the default directory for this file, exactly which Lmhosts file is consulted depends on the value of the **DataBasePath** entry in the following registry subkey:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

The **DataBasePath** entry tells the local computer where to look for the Lmhosts file.



Caution

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

Long Connect Times by Using Lmhosts

To determine the cause of long connect times after adding an entry to Lmhosts, take a look at the order of the entries in the Lmhosts file.

Long connect times can occur when a large Lmhosts file has the specified entry at the end of the file. To speed up resolution of the entry, mark the entry in Lmhosts as a preloaded entry by following the mapping with the #PRE tag. Then use the **nbtstat -R** command to update the local NetBIOS name cache immediately.

Alternately, you can place the mapping higher in the Lmhosts file. The Lmhosts file is parsed sequentially from the top to locate entries without the #PRE keyword. Therefore, you should always place frequently used entries near the top of the file and place the #PRE entries near the bottom. For more information about mapping an entry in the Lmhosts file by using the #PRE keyword, see *Microsoft Windows 2000 Server TCP/IP Core Networking Guide*.

Check the WINS Configuration

Make sure your computer's WINS configuration is correct. In particular, check the address for the WINS server.

► To examine your WINS configuration

1. In Control Panel, double-click **Network Connections**.
2. Right-click the connection you want to examine, and then click **Properties**.
3. On the **General Tab**, select **Internet Protocol (TCP/IP)**, and then click **Properties**.
4. In the **Internet Protocol (TCP/IP) Properties** dialog box, click **Advanced**.
5. In the **Advanced TCP/IP Settings** dialog box, click the **WINS** tab.

In the **WINS Configuration** dialog box, add the IP address of the server (if none is listed) and check to see whether Lmhosts lookup is enabled. Also check to see whether NetBIOS

over TCP/IP is taken from the DHCP server, enabled, or disabled. If you are using DHCP for this host computer, take the value from the DHCP server; otherwise, enable NetBIOS over TCP/IP.

Troubleshooting IP Routing

Windows Server 2003 supports IP routing on both single-homed and multihomed computers with or without the Routing and Remote Access service. Routing and Remote Access includes the Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) routing protocols. Routers can use RIP or OSPF to dynamically exchange routing information.

This section provides information about the Windows Server 2003–based routing table as used on singlehomed and multihomed computers with or without Routing and Remote Access. This background information helps with TCP/IP troubleshooting. For more information about TCP/IP unicast and multicast routing, see *Microsoft Windows 2000 Server Resource Kit Internetworking Guide*.

Cannot Connect to a Specific Server

If you are having problems trying to connect to a specific server by using NetBIOS-based connections, use the **nbtstat -n** command to determine what NetBIOS names the server used to register on the network.

Nbtstat -n output lists several names that the computer has registered. A name resembling the computer's name as shown on the desktop should be present. If not, try one of the other unique names displayed by Nbtstat.

The Nbtstat tool can also display the cached entries from either PRE entries in the Lmhosts file or from recently resolved NetBIOS names. If the NetBIOS name that the computers are using for the server is the same, and the other computers are on a remote subnet, be sure that they have the computer's mapping in their Lmhosts files or WINS servers.

Multiple-Use Servers

Servers with other tasks can be providing routing services, particularly in LANs. Provided you can access a server you have found to be problematic, especially intermittently, you can verify that no other services or applications are tying up server resources, or that someone is not loading large files on to or off of the server. This applies to the client computer as well. One method to verify this would be to open Computer Management by right-clicking My Computer and selecting Manage, expanding Shared Folders and selecting the Sessions folder to show who has a current session with your computer. From the Computer Management window, you can also open Event Viewer, which can show who accessed the server, or if other problems are occurring at certain times or with a specific frequency. Another method would be to open up Task Manager to verify which applications or processes are running, and how they are affecting the CPU usage.

Hanging Connection to Remote Host

To determine why a TCP/IP connection to a remote computer is not working properly, use the **netstat -a** command to show the status of all activity on TCP and UDP ports on the local computer.

A good TCP connection usually shows 0 bytes in the Sent and Received queues. If data is blocked in either queue, the connection is probably faulty. If not, you are probably experiencing network or application delay.

Examining the Routing Table with Route

In order for two hosts to exchange IP datagrams, they must both have a route to each other, or use default gateways that know of a route. Normally, routers exchange information with each other by using a routing protocol such as RIP.

Enabling IP Routing

By default, IP routing is disabled. To enable IP routing, you must allow the computer to forward IP packets it receives. If you are not using the Routing and Remote Access service to enable routing, you must manually modify the registry.



Caution

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

► To enable IP routing

1. Click **Start**, click **Run**, type **regedit.exe**, and then click **OK**.
2. In the registry editor, navigate to **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters**
3. Select the **IPEnableRouter** entry.
4. Click **Edit**, and then click **Modify**.
5. Under **Value data**, set the value of this field to **1**, to enable IP routing for all network connections installed and used by this computer.
6. Close the registry editor.

If a Windows Server 2003 router does not have an interface on a given subnet, it needs a route to get to that subnet. This can be handled by using a default route or by adding static routes. To add a static route, use either Routing and Remote Access or the **route add** command. For example:

```
route add 172.16.41.0 mask 255.255.255.0 172.16.40.1 metric 2
```

In this example, the **route add** command states that to reach the 172.16.41.0 subnet with a mask of 255.255.255.0, use gateway 172.16.40.1. It also shows that the subnet is two hops away. You

might need to add static routes on downstream routers indicating how to get back to the 172.16.40.0/24 subnet.

If you want the route to persist after a reboot, add the **-p** option to the end of the command line. Persistent routes are recorded in the registry.

Examining Paths with Tracert

Tracert is a route tracing utility that uses incrementally higher values in the TTL field in the IP header to determine the route from one host to another through a network. It does this by sending ICMP Echo messages and analyzing ICMP error messages that return. Tracert allows you to determine the path of a forwarded packet from router to router for up to 30 hops. If a router has failed or if the packet is routed into a loop, Tracert reveals the problem. After the problem router is found, its administrator can be contacted if it is an offsite router, or the router can be restored to fully functional status if it is under your control.

Troubleshooting Gateways

If you see the message “Your default gateway does not belong to one of the configured interfaces...” during the installation and setup of the operating system, find out whether the default gateway is located on the same logical network as the computer’s network adapter. The easiest way to do this is to compare the network ID portion of the default gateway’s IP address with the network IDs of the computer’s network adapters. In other words, check that the bitwise logical AND of the IP address and the subnet mask equals the bitwise logical AND of the default gateway and the subnet mask.

For example, a computer with a single network adapter configured with an IP address of 172.16.27.139 and a subnet mask of 255.255.0.0 requires a default gateway of the form 172.16.y.z. The network ID of the IP interface is 172.16.0.0/16. Using the subnet mask, TCP/IP can determine that all traffic on this network is local; everything else must be sent to the gateway.

Troubleshooting Proxy ARP

Proxy ARP is the answering of ARP requests on behalf of another node. As described in RFC 925, proxy ARP is used in situations in which a subnet is divided without the use of a router. A proxy ARP device is placed between nodes that are not on the same logical subnet. The proxy ARP answers ARP requests and facilitates the forwarding of unicast IP packets for communication between nodes on separate segments.

Examples of proxy ARP devices include:

- Routing and Remote Access in computers running Windows Server 2003, which uses proxy ARP to facilitate communications between remote access clients and nodes on the network segment to which the remote access server is attached.
- The Network Bridge feature in Windows XP and Windows Server 2003 Standard Edition and the 32-bit version of Windows Server 2003 Enterprise Edition, which acts as a proxy ARP device when performing Layer 3 bridging between network segments.

Network traffic sometimes fails because a router's proxy ARP request returns the wrong address. A router makes this ARP request on behalf of an IP address on its internal subnets (just as a remote access server makes a request on the LAN for its remote access clients). The problem is that the router's proxy ARP requests return the wrong MAC address to the sending host. As a result, the sending host sends its traffic to the wrong MAC address. In other words, the problem stems from proxy ARP replies.

To address this problem, use Network Monitor to capture a trace. If the trace reveals that when the sending host sends an ARP request for the MAC address of the destination IP address and a device (usually a router) replies with an incorrect MAC address, the IP address-to-MAC address resolution on the destination computer might be incorrect. For more information about Network Monitor, on your computer running Windows Server 2003 click **Start**, click **Run**, type **hh NETMNconcepts.chm**, and then click **OK**.

To determine if this is the problem, check the ARP cache of the source host to make sure it is getting the correct IP address to MAC address resolution. Alternatively, you can capture all traffic with Network Monitor and later filter the captured traffic to display only the ARP and Reverse Address Resolution Protocol (RARP) protocols. The RARP protocol converts MAC addresses to IP addresses and is defined in RFC 903.

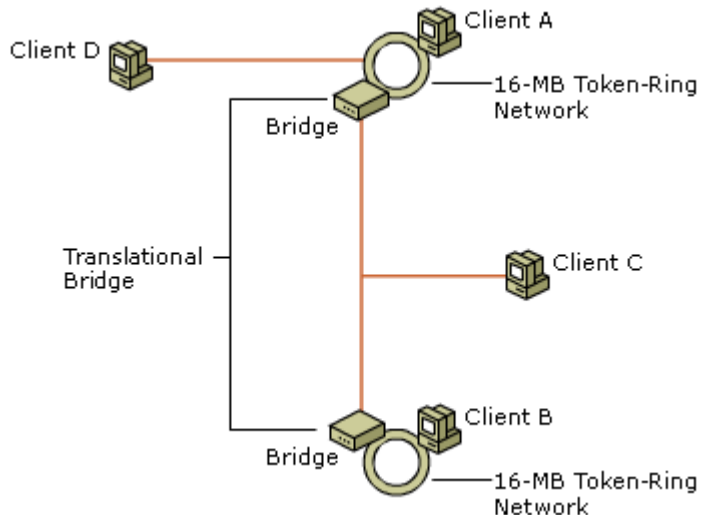
You can fix the ARP problem by disabling 'Proxy ARP' on the offending device. Exactly how this is done depends on the device's make and model; consult the manufacturer's documentation.

Troubleshooting Translational Bridging

Allowing nodes from Ethernet segments to communicate with nodes on Token Ring rings presents a number of challenges. Following are the most common problems associated with translational bridging.

The primary factor responsible for problems in this situation are differing Maximum Transmission Units (MTUs) between segments. Token Ring MTUs range from 4,464 to 17,914 bytes; the Ethernet MTU is 1,500. A FDDI segment has an MTU of 4,532 bytes. When a bridge or Layer 2 switch connects two of these differing networking technologies, packets can be dropped because the Layer 2 switch cannot fragment the data and cannot alert the sending node of the reduced MTU.

In the example shown in Figure 12.1, an Ethernet backbone connects two 16-MB token rings. Instead of a router, a translational bridge in the form of a Layer 2 switch connects the segments. In this case, local traffic on the Token Rings uses an MTU of 17,914 and is not affected by the bridge. However, when Client A must communicate with Client B, the bridge drops large packets without notifying Client A of the need to fragment. In this situation, Client A has no way to discover the MTU on the other side of the bridge.

Figure 12.1 Connecting Two Token Ring Networks with an Ethernet Bridge

Other symptoms of translational bridging problems might include the ability to ping a computer on the far side of the bridge, being able to establish a connection, but not being able to send bulk data. This occurs because Echo messages and TCP connection establishment segments are small. When sending bulk data, however, large segments at the size of the MTU of the locally attached network are sent and dropped by the Layer 2 switch. Another example is when a computer is able to use FTP to establish a session, but is unable to use a `get FileName` command, which requires sending a larger packet over the switch.

In Windows Server 2003, the **MTU** registry entry can be adjusted to meet the MTU requirement of the Ethernet segment connecting the two Token Ring segments, reducing all MTUs to the lowest MTU supported on all the links of the subnet. Each node's MTU is reduced to 1,500 bytes to meet the requirements of the Ethernet backbone. However, this solution requires that all traffic (even traffic that is local on a Token Ring) is sent within the reduced MTU.

Using Ping to Determine Maximum Transmission Units

You can use the **ping -l -f** command to send packets with a defined ICMP Echo data size, by using the **-f** option to specify that the packets will not be fragmented. By sending packets of varying sizes, you can determine the MTU for any given bridge by noting which packet sizes cross the bridge successfully. For example, in Figure 12.1, a ping packet can be sent from Client A to Client C with a size of 1,472 bytes, which generates an Echo Reply packet from Client C. However, if a size of 1,473 bytes or greater is used, the intermediate switch drops the packet. Client C does not receive the Echo and no Echo Reply is generated.

The default ICMP Echo contains 32 bytes of data; you can use the **ping IPAddress** or **ping HostName -l DataSize** command to specify a different data size. For example, you can ping with the maximum Ethernet data size by entering this command:

```
ping 134.56.78.1. -l 1472
```

The data size specified by the **-I** option is 1,472 rather than the Ethernet IP MTU of 1,500 because 20 bytes are reserved to make room for the IP header and 8 bytes must be allocated for the ICMP Echo header.

When you have determined the MTU, you can set the packet size on either side of the bridge by changing the value of the registry entry. The **MTU** registry entry can be found in the following subkey:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\Adapter_GUID

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

For more information about MTU, see the Microsoft Press book *Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference*.

Troubleshooting PMTU Black-Hole Routers

Some routers do not send an “ICMP Destination Unreachable-Fragmentation Needed and DF Set” message when they cannot forward an IP datagram. Instead, they silently discard the datagram. Typically, an IP datagram cannot be forwarded because its maximum segment size is too large for the receiving server, and the Don’t Fragment bit is set in the header of the datagram. Routers that ignore these datagrams and send no message are called PMTU “black-hole” routers.

To respond effectively to black-hole routers, you must enable the PMTU black-hole detection feature of TCP/IP. PMTU black-hole detection Detect recognizes repeated unacknowledged transmissions and responds by turning off the Don’t Fragment bit. After a datagram is transmitted successfully, it reduces the maximum segment size and turns the Don’t Fragment bit on again.

The PMTU black-hole detection feature is disabled by default, but you can enable it by adding the **EnablePMTUBHDetect** entry to the registry and setting its value to **1**.

EnablePMTUBHDetect is an optional entry that does not appear in the registry unless you add it to the following subkey:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

**Caution**

Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on the computer.

You can disable PMTU black-hole detection by deleting **EnablePMTUBHDetect** from the registry or by setting the entry’s value to **0**.

A second registry entry, **EnablePMTUDiscovery**, also helps address the PMTU black-hole router problem. This key is enabled by default. **EnablePMTUDiscovery** completely enables or disables the PMTU discovery mechanism. When PMTU discovery is disabled, a TCP Maximum Segment Size (MSS) of 536 bytes is used for all non-local destination addresses.

Discovering PMTU with Ping

The PMTU between two hosts can be discovered manually by using the **ping -f** command, as follows:

```
ping -f -n NumberOfPings [-l size] DestinationIPAddress
```

The following example shows how Ping's size parameter can be varied until the MTU is found. Note that Ping's size parameter specifies just the size of the ICMP Echo data to send, not including the IP and ICMP Echo headers. The ICMP Echo header is 8 bytes, and the IP header is normally 20 bytes. In the Ethernet case shown here, the link layer MTU contains the maximum-sized Ping buffer plus 28, for a total of 1500 bytes on the first ping and 1501 on the second:

```
C:\>ping -f -n 1 -l 1472 10.99.99.10
Pinging 10.99.99.10 with 1472 bytes of data:
Reply from 10.99.99.10: bytes=1472 time<10ms TTL=128
Ping statistics for 10.99.99.10:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\>ping -f -n 1 -l 1473 10.99.99.10
Pinging 10.99.99.10 with 1473 bytes of data:
Packet needs to be fragmented but DF set.
Ping statistics for 10.99.99.10:
    Packets: Sent = 1, Received = 0, Lost = 1 (100% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

In the second ping, the IP layer returns an ICMP error message that Ping interprets. If the router had been a black-hole router, Ping would not be answered after its size exceeded the MTU that the router could handle. Ping can be used in this manner to detect such a router.

Troubleshooting Services

In addition to its role in providing basic network communications, TCP/IP is the cornerstone of a number of network services such as Routing and Remote Access, printing, IPSec, and the Computer Browser service. These services are discussed in more detail in other chapters, but a few examples of basic troubleshooting for these services are described below.

Cannot Ping Across a Router as a Remote Access Client

This problem occurs if you have selected **Use default gateway on remote network** under the **General** tab of the **Advanced Internet Protocol (TCP/IP) Properties** in the **Dial-Up Connections** page. This feature adds a default route to the routing table with a lower metric than the existing default route, adjusting the metric of the existing default route if needed. All non-local traffic is now forwarded to the gateway on the remote access link. However, to access the Internet, this feature must be enabled.

To ping or otherwise connect to computers in a remote subnet across a router while you are connected as a remote access client to a remote Windows remote access server, use the **route add** command to add the route of the subnet you want to use.

Troubleshooting TCP/IP Database Files

Table 12.2 lists the UNIX-style database files that are stored in the `systemroot\System32\Drivers\Etc` directory when you install Microsoft TCP/IP:

Table 12.2 TCP/IP Database File

File Name	Use
Hosts	Provides host name-to-IP-address resolution for Windows Sockets applications.
Lmhosts.sam	Sample file for Lmhosts, which provides NetBIOS name-to-remote-IP-address resolution for NetBIOS applications.
Networks	Provides network name-to-network ID resolution for Windows Sockets applications.
Protocols	Provides protocol name-to-protocol ID resolution for Windows Sockets applications.
Services	Provides service name-to-port ID resolution for Windows Sockets applications.

To troubleshoot any of these files on a local computer, make sure the entry format in each file matches the format defined in the sample file originally installed with Microsoft TCP/IP. Check for spelling errors, invalid IP addresses, and identifiers.

Additional Resources

- For more information about TCP/IP protocols and processes and Windows implementation details, see *Microsoft Windows Server 2003 TCP/IP Protocols and Services Technical Reference* by Joseph Davies and Thomas Lee, 2003, Redmond, Washington: Microsoft Press.
- For more information about TCP/IP, see *Microsoft Windows 2000 Server TCP/IP Core Networking Guide*, by Microsoft Press, 2000, Redmond, Washington: Microsoft Press.
- For more information about TCP/IP troubleshooting, see *Windows NT TCP/IP Network Administration* by Craig Hunt and Robert Bruce Thompson, 1998, Sebastopol, California: O'Reilly.

- For more information about Network Monitor, see *Microsoft Windows Server 2003 Administrator's Companion*.