

# Windows 2003 TCP/IP Technical Reference

## (Microsoft Corporation)

Transmission Control Protocol/Internet Protocol, or TCP/IP, is a protocol suite developed for communications over, often dissimilar, networks. Microsoft Windows Server 2003 provides extensive support for the TCP/IP suite for IP version 4 (IPv4), as both a protocol and a set of services for connectivity and management of IP internetworks. Knowledge of the basic concepts of TCP/IP is an absolute requirement for the proper understanding of the configuration, deployment, and troubleshooting of Windows Server 2003.

This subject describes how TCP/IP with IPv4 relates to other networking protocols, the functions it performs, how addresses are structured and assigned, and how packets are structured and routed.

### What Is TCP/IP?

Transmission Control Protocol/Internet Protocol, or TCP/IP, is an industry-standard suite of protocols designed for large internetworks. TCP/IP, which was developed in 1969 by the U.S. Department of Defense (DoD) Advanced Research Projects Agency (DARPA), is the result of a resource-sharing experiment called Advanced Research Projects Agency Network (ARPANET). The TCP/IP protocol was developed to provide high-speed communication network links. Since 1969, ARPANET has grown into a worldwide community of networks known as the Internet.

Before TCP/IP, there was no way for computers to communicate easily and securely on public networks. Windows Server 2003 TCP/IP was designed to make it easy to integrate Microsoft systems into large-scale corporate, government, and public networks, and to provide the ability to operate over those networks in a secure manner. The Windows Server 2003 TCP/IP protocol is installed by default and, unlike previous versions of Windows, cannot be uninstalled. However, you can reset the TCP/IP configuration to a default state with the netsh interface ip reset command.

The Windows TCP/IP suite contains core protocol elements, services, and the interfaces between them. The Transport Driver Interface (TDI) and the Network Device Interface Specification (NDIS) are public, and their specifications are available from Microsoft. In addition, there are a number of higher-level interfaces available to user-mode applications. The most commonly used are Windows Sockets and NetBIOS.

### Windows Server 2003 TCP/IP

Windows Server 2003 TCP/IP enables enterprise networking and connectivity. Adding TCP/IP to a Windows Server 2003 configuration offers the following advantages:

- A standard, routable enterprise networking protocol that is the most complete and accepted protocol available. All modern network operating systems offer TCP/IP support, and most large networks rely on TCP/IP for much of their network traffic.
- A technology for connecting dissimilar systems. Many standard connectivity utilities are available to access and transfer data between dissimilar systems, including File Transfer Protocol (FTP) and Telnet, a terminal emulation protocol. Several of these standard utilities are included with Windows Server 2003.
- A robust, scalable, cross-platform client/server framework. Windows Server 2003 TCP/IP offers the Windows Sockets interface, which is ideal for developing client/server applications that can run on Windows Sockets-compliant TCP/IP protocol implementations from other vendors.
- A method of gaining access to the Internet. The Internet consists of thousands of networks worldwide, connecting research facilities, universities, libraries, private companies, and individuals.

# Windows 2003 TCP/IP Technical Reference

## (Microsoft Corporation)

**Note:** The term Internet refers to the worldwide public Internet. An intranet refers to a private IP-based internetwork.

### Support for Standard Features

Windows Server 2003 TCP/IP supports the following standard features:

- Ability to bind to multiple network adapters with different media types.
- Logical and physical multihoming.
- Internal IP routing capability.
- Internet Group Management Protocol (IGMP) (IP multicasting).
- Duplicate IP address detection.
- Multiple configurable default gateways.
- Dead gateway detection.
- Automatic Path Maximum Transmission Unit (PMTU) discovery.
- Internet Protocol security (IPSec).
- Quality of Service (QoS).
- Asynchronous Transfer Mode Tutorial (ATM) services.
- Virtual Private Networks (VPNs) with the Point-to-Point Tunneling Protocol (PPTP) and the Layer Two Tunneling Protocol with IPSec (L2TP/IPSec).

### Windows Server 2003 TCP/IP Features

The features and improvements of TCP/IP for Windows Server 2003 include the following:

- Windows Server 2003 and Windows XP (Service Pack 1 or later) now include a production-quality IPv6 protocol stack. For more information about IPv6, see IPv6 Technical Reference.
- Auto-negotiation of RFC 1323 options (window scaling and TCP timestamps).
- Default support of network interface cards providing large send offload (LSO) and checksum offload.
- Internet Group Management Protocol (IGMP) version 3.
- Reliable multicast with Pragmatic General Multicast (PGM).
- Alternate configuration of a static IP address configuration.
- Automatic determination of the interface-related and default route metrics.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## TCP/IP Standards

The standards for TCP/IP are published in a series of documents that are called Requests for Comments (RFCs). RFCs describe the internal workings of the Internet. Some RFCs describe network services or protocols and their implementations, whereas others summarize policies. TCP/IP standards are always published as RFCs, although not all RFCs specify standards.

TCP/IP standards are not developed by a committee, but rather by consensus. Anyone can submit a document for publication as an RFC. Documents are reviewed by a technical expert, a task force, or the RFC editor, and then assigned a status. The status specifies whether a document is being considered for becoming a standard.

There are five RFC status assignments, shown below:

### RFC Status Assignments

Status	Description
Required	Must be implemented on all TCP/IP-based hosts and gateways.
Recommended	Encouraged that all TCP/IP-based hosts and gateways implement the RFC specifications. Recommended RFCs are usually implemented.
Elective	Implementation is optional. Application has been agreed to but is not a requirement.
Limited Use	Not intended for general use.
Not Recommended	Not recommended for implementation.

If a document is being considered for becoming a standard, it goes through stages of development, testing, and acceptance known as the Internet Standards Process. These stages are formally labeled maturity levels. The following table lists the three maturity levels for Internet Standards.

### Maturity Levels for Internet Standards

Maturity Level	Description
Proposed Standard	This specification is generally stable, has resolved known design issues, is believed to be well understood, has received significant community review, and appears to hold enough community interest to be considered valuable.
Draft Standard	This must be well understood and known to be quite stable, both in its semantics and as a basis for developing an implementation.
Internet Standard	This specification is characterized by a high degree of technical maturity, and it is generally agreed that the specified protocol or service provides significant benefit to the Internet community.

When a document is published, it is assigned an RFC number. If changes are required, a new RFC is published with a new number. The original RFC is never updated. Therefore, it is important to verify that you have the most recent RFC on a particular topic.

RFCs can be obtained in several ways. To obtain any RFC, or a full and current indexed listing of all RFCs published to date, see the IETF RFC Database.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## TCP/IP Core Protocols

The TCP/IP protocol component that is installed in your network operating system is a series of interconnected protocols called the TCP/IP core protocols. All other applications and other protocols in the TCP/IP protocol suite rely on the basic services provided by the following protocols: IP, ARP, ICMP, IGMP, TCP, and UDP. For more information about these protocols, see How TCP/IP Works.

## TCP/IP Services

The Windows Server 2003 operating system provides the following TCP/IP-related services:

- Dynamic Host Configuration Protocol (DHCP) client and server and DHCP Relay Agent (with the Routing and Remote Access service).
- In the absence of a DHCP server, Automatic Private IP Addressing (APIPA) or an alternate static IP address configuration is used.
- Windows Internet Name Service (WINS), a NetBIOS name client and server.
- Domain Name System (DNS) client and server, including support for DNS dynamic updates.
- Dial-up support using the Point-to-Point Protocol (client and server) and Serial Line Internet Protocol (client only).
- PPTP and L2TP/IPSec, used for remote access and site-to-site VPN connections.
- TCP/IP network printing (client only with the Lpr.exe and Lpq.exe tools).
- SNMP agent.
- NetBIOS interface.
- Network Location Service.
- Windows Sockets version 2 (Winsock2) interface.
- Remote Procedure Call (RPC) support.
- Network Dynamic Data Exchange (NetDDE).
- Computer browsing (My Network Places) across IP routers.
- Reliable multicast with the Pragmatic General Multicast (PGM) protocol.
- Basic TCP/IP connectivity utilities, including: finger, ftp, rcp, rexec, rsh, telnet, and tftp.
- Server and client software for simple network protocols, including: Character Generator, Daytime, Discard, Echo, and Quote of the Day.
- Routing Information Protocol (RIP) listener (for Windows XP Professional) and RIP and Open Shortest Path First (OSPF) (with the Routing and Remote Access service).
- Network Address Translator (NAT) capabilities using either the Internet Connection Service (ICS) or the NAT/Basic Firewall routing protocol component of the Routing and Remote Access Service.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

- Stateful firewalling capabilities using either the Internet Connection Firewall (ICF) or the NAT/Basic Firewall routing protocol component of the Routing and Remote Access service.
- Multicast forwarding and IGMP router and proxy capabilities with the Routing and Remote Access Service.
- TCP/IP management and diagnostic tools, including: arp, ipconfig, nbtstat, netsh, netstat, ping, pathping, route, nslookup, and tracert.

## How TCP/IP Works

TCP/IP for IP version 4 (IPv4) is a networking protocol suite that Microsoft Windows uses to communicate over the internet with other computers. It interacts with Windows naming services like DNS and security technologies, such as IPsec primarily, as these help facilitate the successful and secure transfer of IP packets between machines.

Ideally, TCP/IP is used whenever Windows-based computers communicate over networks.

This subject describes the components of the TCP/IP Protocol Suite, the protocol architecture, which functions TCP/IP performs, how addresses are structured and assigned, and how packets are structured and routed.

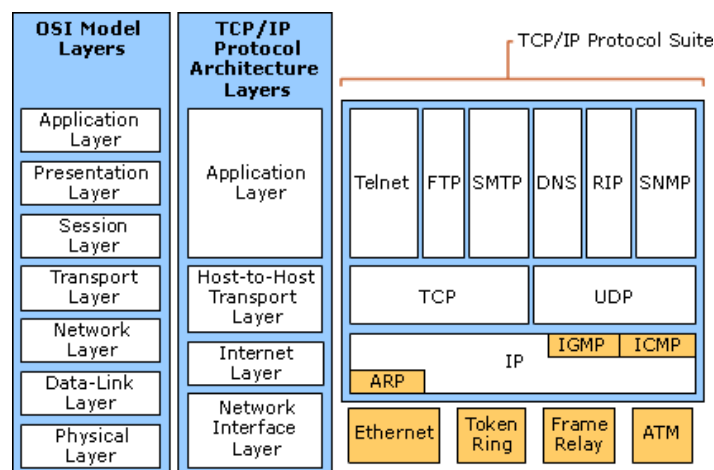
Microsoft Windows Server 2003 provides extensive support for the Transmission Control Protocol/Internet Protocol (TCP/IP) suite, as both a protocol and a set of services for connectivity and management of IP internetworks. Knowledge of the basic concepts of TCP/IP is an absolute requirement for the proper understanding of the configuration, deployment, and troubleshooting of IP-based Windows Server 2003 and Microsoft Windows 2000 intranets.

## TCP/IP Protocol Architecture

TCP/IP protocols map to a four-layer conceptual model known as the DARPA model, named after the U.S. government agency that initially developed TCP/IP. The four layers of the DARPA model are: Application, Transport, Internet, and Network Interface. Each layer in the DARPA model corresponds to one or more layers of the seven-layer Open Systems Interconnection (OSI) model.

The following figure shows the TCP/IP protocol architecture:

**TCP/IP Protocol Architecture**



# Windows 2003 TCP/IP Technical Reference

## (Microsoft Corporation)

Note: The architectural diagram above corresponds to the Internet Protocol TCP/IP and does not reflect IP version 6. To see a TCP/IP architectural diagram that includes IPv6, see How IPv6 Works in this technical reference.

### Network Interface Layer

The Network Interface layer (also called the Network Access layer) handles placing TCP/IP packets on the network medium and receiving TCP/IP packets off the network medium. TCP/IP was designed to be independent of the network access method, frame format, and medium. In this way, TCP/IP can be used to connect differing network types. These include local area network (LAN) media such as Ethernet and Token Ring and WAN technologies such as X.25 and Frame Relay. Independence from any specific network media allows TCP/IP to be adapted to new media such as asynchronous transfer mode (ATM).

The Network Interface layer encompasses the Data Link and Physical layers of the OSI model. Note that the Internet layer does not take advantage of sequencing and acknowledgment services that might be present in the Network Interface layer. An unreliable Network Interface layer is assumed, and reliable communication through session establishment and the sequencing and acknowledgment of packets is the function of the Transport layer.

### Internet Layer

The Internet layer handles addressing, packaging, and routing functions. The core protocols of the Internet layer are IP, ARP, ICMP, and IGMP.

- The Internet Protocol (IP) is a routable protocol that handles IP addressing, routing, and the fragmentation and reassembly of packets.
- The Address Resolution Protocol (ARP) handles resolution of an Internet layer address to a Network Interface layer address, such as a hardware address.
- The Internet Control Message Protocol (ICMP) handles providing diagnostic functions and reporting errors due to the unsuccessful delivery of IP packets.
- The Internet Group Management Protocol (IGMP) handles management of IP multicast group membership.

The Internet layer is analogous to the Network layer of the OSI model.

### Transport Layer

The Transport layer (also known as the Host-to-Host Transport layer) handles providing the Application layer with session and datagram communication services. The core protocols of the Transport layer are Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP).

- TCP provides a one-to-one, connection-oriented, reliable communications service. TCP handles the establishment of a TCP connection, the sequencing and acknowledgment of packets sent, and the recovery of packets lost during transmission.
- UDP provides a one-to-one or one-to-many, connectionless, unreliable communications service. UDP is used when the amount of data to be transferred is small (such as data that fits into a single packet), when you do not want the overhead of establishing a TCP connection, or when the applications or upper layer protocols provide reliable delivery.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

The TCP/IP Transport layer encompasses the responsibilities of the OSI Transport layer.

### Application Layer

The Application layer lets applications access the services of the other layers and defines the protocols that applications use to exchange data. There are many Application layer protocols and new protocols are always being developed.

The most widely known Application layer protocols are those used for the exchange of user information:

- The Hypertext Transfer Protocol (HTTP) is used to transfer files that make up the Web pages of the World Wide Web.
- The File Transfer Protocol (FTP) is used for interactive file transfer.
- The Simple Mail Transfer Protocol (SMTP) is used for the transfer of mail messages and attachments.
- Telnet, a terminal emulation protocol, is used for logging on remotely to network hosts.

Additionally, the following Application layer protocols help facilitate the use and management of TCP/IP networks:

- The Domain Name System (DNS) is used to resolve a host name to an IP address.
- The Routing Information Protocol (RIP) is a routing protocol that routers use to exchange routing information on an IP internetwork.
- The Simple Network Management Protocol (SNMP) is used between a network management console and network devices (routers, bridges, intelligent hubs) to collect and exchange network management information.

Examples of Application layer interfaces for TCP/IP applications are Windows Sockets and NetBIOS. Windows Sockets provides a standard application programming interface (API) under Windows Server 2003. NetBIOS is an industry-standard interface for accessing protocol services such as sessions, datagrams, and name resolution. More information on Windows Sockets and NetBIOS is provided later in this chapter.

The TCP/IP Application layer encompasses the responsibilities of the OSI Session, Presentation, and Application layers.

### TCP/IP Core Protocols

The TCP/IP protocol component that is installed in your network operating system is a series of interconnected protocols called the core protocols of TCP/IP. All other applications and other protocols in the TCP/IP protocol suite rely on the basic services provided by the following protocols: IP, ARP, ICMP, IGMP, TCP, and UDP.

### IP

IP is a connectionless, unreliable datagram protocol primarily responsible for addressing and routing packets between hosts. Connectionless means that a session is not established before exchanging data. Unreliable means that delivery is not guaranteed. IP always makes a “best effort” attempt to deliver a packet. An IP packet might be lost, delivered out of sequence, duplicated, or delayed. IP

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

does not attempt to recover from these types of errors. The acknowledgment of packets delivered and the recovery of lost packets is the responsibility of a higher-layer protocol, such as TCP. IP is defined in RFC 791.

An IP packet consists of an IP header and an IP payload. The following table describes the key fields in the IP header.

**Key Fields in the IP Header**

IP Header Field	Function
Source Address	The IP address of the original source of the IP datagram.
Destination Address	The IP address of the final destination of the IP datagram.
Identification	Used to identify a specific IP datagram and to identify all fragments of a specific IP datagram if fragmentation occurs.
Protocol	Informs IP at the destination host whether to pass the packet up to TCP, UDP, ICMP, or other protocols.
Checksum	A simple mathematical computation used to verify the bit-level integrity of the IP header.
Time to Live (TTL)	Designates the number of network segments on which the datagram is allowed to travel before being discarded by a router. The TTL is set by the sending host and is used to prevent packets from endlessly circulating on an IP internetwork. When forwarding an IP packet, routers are required to decrease the TTL by at least one.

### Fragmentation and Reassembly

If a router receives an IP packet that is too large for the network to which the packet is being forwarded, IP fragments the original packet into smaller packets that fit on the downstream network. When the packets arrive at their final destination, IP on the destination host reassembles the fragments into the original payload. This process is referred to as fragmentation and reassembly. Fragmentation can occur in environments that have a mix of networking media, such as Ethernet and Token Ring.

The fragmentation and reassembly works as follows:

- When an IP packet is sent by the source, it places a unique value in the Identification field.
- The IP packet is received at the router. The IP router notes that the maximum transmission unit (MTU) of the network onto which the packet is to be forwarded is smaller than the size of the IP packet.
- IP divides the original IP payload into fragments that fit on the next network. Each fragment is sent with its own IP header that contains:
  - The original Identification field identifying all fragments that belong together.
  - The More Fragments Flag indicating that other fragments follow. The More Fragments Flag is not set on the last fragment, because no other fragments follow it.
  - The Fragment Offset field indicating the position of the fragment relative to the original IP payload.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

When the fragments are received by IP at the remote host, they are identified by the Identification field as belonging together. The Fragment Offset field is then used to reassemble the fragments into the original IP payload.

### ARP

When IP packets are sent on shared access, broadcast-based networking media — such as Ethernet or Token Ring — the media access control (MAC) address corresponding to a forwarding IP address must be resolved. ARP uses MAC-level broadcasts to resolve a known forwarding or next-hop IP address to its MAC address. ARP is defined in RFC 826.

### ICMP

Internet Control Message Protocol (ICMP) provides troubleshooting facilities and error reporting for packets that are undeliverable. For example, if IP is unable to deliver a packet to the destination host, ICMP sends a Destination Unreachable message to the source host. The following table shows the most common ICMP messages.

#### Common ICMP Messages

ICMP Message	Function
Echo Request	Troubleshooting message used to check IP connectivity to a desired host. The ping utility sends ICMP Echo Request messages.
Echo Reply	Response to an ICMP Echo Request.
Redirect	Sent by a router to inform a sending host of a better route to a destination IP address.
Source Quench	Sent by a router to inform a sending host that its IP datagrams are being dropped due to congestion at the router. The sending host then lowers its transmission rate. Source Quench is an elective ICMP message and is not commonly implemented.
Destination Unreachable	Sent by a router or the destination host to inform the sending host that the datagram cannot be delivered.

The following table describes the most common ICMP Destination Unreachable ICMP messages.

#### Common ICMP Destination Unreachable Messages

Destination Unreachable Message	Description
Host Unreachable	Sent by an IP router when a route to the destination IP address cannot be found.
Protocol Unreachable	Sent by the destination IP node when the <b>Protocol</b> field in the IP header cannot be matched with an IP client protocol currently loaded.
Port Unreachable	Sent by the destination IP node when the Destination Port in the UDP header cannot be matched with a process using that port.
Fragmentation Needed and DF Set	Sent by an IP router when fragmentation must occur but is not allowed due to the source node setting the <b>Don't Fragment (DF)</b> flag in the IP header.
Source Route Failed	Sent by an IP router when delivery of the IP packet using source route information (stored as source route option headers) fails.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

ICMP does not make IP a reliable protocol. ICMP attempts to report errors and provide feedback on specific conditions. ICMP messages are carried as unacknowledged IP datagrams and are themselves unreliable. ICMP is defined in RFC 792.

### IGMP

Internet Group Management Protocol (IGMP) is a protocol that manages host membership in IP multicast groups on a network segment. An IP multicast group, also known as a host group, is a set of hosts that listen for IP traffic destined for a specific IP multicast address. IP multicast traffic is sent to a single MAC address but processed by multiple IP hosts. A specific host listens on a specific IP multicast address and receives all packets to that IP address.

The following are some of the additional aspects of IP multicasting:

- Host group membership is dynamic, hosts can join and leave the group at any time.
- A host group can be of any size.
- Members of a host group can span IP routers across multiple networks. This situation requires IP multicast support on the IP routers and the ability for hosts to register their group membership with local routers. Host registration is accomplished using IGMP.
- A host can send traffic to an IP multicast address without belonging to the corresponding host group.

For a host to receive IP multicasts, an application must inform IP that it will receive multicasts at a specified IP multicast address. If the network technology supports hardware-based multicasting, the network interface is told to pass up packets for a specific IP multicast address. In the case of Ethernet, the network adapter is programmed to respond to a multicast MAC address corresponding to the specified IP multicast address.

A host supports IP multicast at one of the following levels:

- Level 0: No support to send or receive IP multicast traffic.
- Level 1: Support exists to send but not receive IP multicast traffic.
- Level 2: Support exists to both send and receive IP multicast traffic. Windows Server 2003, Windows 2000, Microsoft Windows NT version 3.5 and later, and TCP/IP support level 2 IP multicasting.

The protocol to register host group information is IGMP, which is required on all hosts that support level 2 IP multicasting. IGMP packets are sent using an IP header.

IGMP messages take three forms:

- **Host Membership Report.** When a host joins a host group, it sends an IGMP Host Membership Report message to the all-hosts IP multicast address (224.0.0.1) or to the specified IP multicast address declaring its membership in a specific host group by referencing the IP multicast address. A host can also specify the specific sources from which multicast traffic is needed.
- **Host Membership Query.** When a router polls a network to ensure that there are members of a specific host group, it sends an IGMP Host Membership Query message to the all-hosts IP

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

multicast address. If no responses to the poll are received after several polls, the router assumes no membership in that group for that network and stops advertising that multicast group information to other routers.

- **Group Leave.** When a host is no longer interested in receiving multicast traffic sent to a specific IP multicast address and it sent the last IGMP Host Membership Report message in response to an IGMP Host Membership Query, it sends an IGMP Group Leave message to the specific IP multicast address. Local routers verify that the host sending the IGMP Group Leave message is the last group member for that multicast address on that subnet. If no responses to the poll are received after several polls, the router assumes no membership in that group for that subnet and stops advertising that multicast group information to other routers.

For IP multicasting to span routers across an internetwork, multicast routing protocols are used by routers to communicate host group information so that each router supporting multicast forwarding is aware of which networks contain members of which host groups. IGMP is defined in RFCs 1112 and 2236.

### TCP

TCP is a reliable, connection-oriented delivery service. The data is transmitted in segments. Connection-oriented means that a connection must be established before hosts can exchange data. Reliability is achieved by assigning a sequence number to each segment transmitted. An acknowledgment is used to verify that the data is received. For each segment sent, the receiving host must return an acknowledgment (ACK) within a specified period for bytes received. If an ACK is not received, the data is retransmitted. TCP is defined in RFC 793.

TCP uses byte-stream communications, wherein data within the TCP segment is treated as a sequence of bytes with no record or field boundaries. The following table describes the key fields in the TCP header.

**Key Fields in the TCP Header**

Field	Function
Source Port	TCP port of sending host.
Destination Port	TCP port of destination host.
Sequence Number	Sequence number of the first byte of data in the TCP segment.
Acknowledgment Number	Sequence number of the byte the sender expects to receive next from the other side of the connection.
Window	Current size of a TCP buffer on the host sending this TCP segment to store incoming segments.
TCP Checksum	Verifies the bit-level integrity of the TCP header and the TCP data.

### TCP Ports

A TCP port provides a specific location for delivery of TCP segments. Port numbers below 1024 are well-known ports and are assigned by the Internet Assigned Numbers Authority (IANA). The following table lists a few well-known TCP ports.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## Well-Known TCP Ports

TCP Port Number	Description
20	FTP (Data Channel)
21	FTP (Control Channel)
23	Telnet
80	HTTP used for the World Wide Web
139	NetBIOS session service

### TCP three-way Handshake

A TCP connection is initialized through a three-way handshake. The purpose of the three-way handshake is to synchronize the sequence number and acknowledgment numbers of both sides of the connection and exchange TCP window sizes or the use of large window sizes or TCP timestamps. The following steps outline the process:

1. The initiator of the TCP connection, typically a client, sends a TCP segment to the server with an initial Sequence Number for the connection and a window size indicating the size of a buffer on the client to store incoming segments from the server.
2. The responder of the TCP connection, typically a server, sends back a TCP segment containing its chosen initial Sequence Number, an acknowledgment of the client's Sequence Number, and a window size indicating the size of a buffer on the server to store incoming segments from the client.
3. The initiator sends a TCP segment to the server containing an acknowledgment of the server's Sequence Number.

TCP uses a similar handshake process to end a connection. This guarantees that both hosts have finished transmitting and that all data was received.

### UDP

UDP provides a connectionless datagram service that offers unreliable, best-effort delivery of data transmitted in messages. This means that neither the arrival of datagrams nor the correct sequencing of delivered packets is guaranteed. UDP does not recover from lost data through retransmission. UDP is defined in RFC 768. UDP is used by applications that do not require an acknowledgment of receipt of data and that typically transmit small amounts of data at one time. NetBIOS name service, NetBIOS datagram service, and SNMP are examples of services and applications that use UDP. The following table describes the key fields in the UDP header.

### Key Fields in the UDP Header

Field	Function
Source Port	UDP port of sending host.
Destination Port	UDP port of destination host.
UDP Checksum	Verifies the bit-level integrity of the UDP header and the UDP data.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## UDP Ports

To use UDP, an application must supply the IP address and UDP port number of the destination application. A port provides a location for sending messages. A port functions as a multiplexed message queue, meaning that it can receive multiple messages at a time. Each port is identified by a unique number. It is important to note that UDP ports are distinct and separate from TCP ports even though some of them use the same number. The following table lists a few well-known UDP ports.

### Well-Known UDP Ports

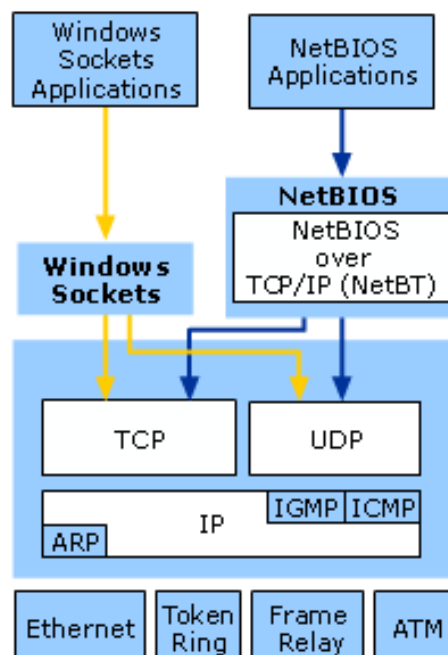
UDP Port Number	Description
53	Domain Name System (DNS) name queries
69	Trivial File Transfer Protocol (TFTP)
137	NetBIOS name service
138	NetBIOS datagram service
161	SNMP

## TCP/IP Application Interfaces

For applications to access the services offered by the core TCP/IP protocols in a standard way, network operating systems like Windows Server 2003 make industry-standard application programming interfaces (APIs) available. APIs are sets of functions and commands that are programmatically called by application code to perform network functions. For example, a Web browser application connecting to a Web site needs access to TCP's connection establishment service.

The following figure shows two common TCP/IP APIs, Windows Sockets and NetBIOS, and their relation to the core protocols.

### APIs for TCP/IP



# Windows 2003 TCP/IP Technical Reference

## (Microsoft Corporation)

### Windows Sockets Interface

The Windows Sockets API is a standard API under Windows Server 2003 for applications that use TCP and UDP. Applications written to the Windows Sockets API run on many versions of TCP/IP. TCP/IP utilities and the SNMP service are examples of applications written to the Windows Sockets interface.

Windows Sockets provides services that allow applications to bind to a particular port and IP address on a host, initiate and accept a connection, send and receive data, and close a connection. There are two types of sockets:

- A stream socket provides a two-way, reliable, sequenced, and unduplicated flow of data using TCP.
- A datagram socket provides a one-way or two-way flow of data using UDP.

A socket is defined by a protocol and an address on the host. The format of the address is specific to each protocol. In TCP/IP, the address is the combination of the IP address and port. Two sockets, one for each end of the connection, form a bi-directional communications path.

To communicate, an application specifies the protocol, the IP address of the destination host, and the port of the destination application. After the application is connected, information can be sent and received.

### NetBIOS Interface

NetBIOS allows applications to communicate over a network. NetBIOS defines two entities, a session-level interface and a session management and data transport protocol.

The NetBIOS interface is a standard API for user applications to submit network input/output (I/O) and control directives to underlying network protocol software. An application program that uses the NetBIOS interface API for network communication can be run on any protocol software that supports the NetBIOS interface.

NetBIOS also defines a protocol that functions at the session/transport level. This is implemented by the underlying protocol software (such as the NetBIOS Frames Protocol NBFP — a component of NetBEUI or NetBIOS over TCP/IP (NetBT)), which performs the network I/O required to accommodate the NetBIOS interface command set. NetBIOS over TCP/IP is defined in RFCs 1001 and 1002. NetBT is enabled by default, however Windows Server 2003 allows you to disable NetBT for an environment that contains no NetBIOS-based network clients or applications.

NetBIOS provides commands and support for NetBIOS Name Management, NetBIOS Datagrams, and NetBIOS Sessions.

### NetBIOS Name Management

NetBIOS name management services provide the following functions:

- Name registration and release. When a TCP/IP host initializes, it registers its NetBIOS names by broadcasting or directing a NetBIOS name registration request to a NetBIOS Name Server such as a WINS server. If another host has registered the same NetBIOS name, either the host or a NetBIOS Name Server responds with a negative name registration response. The initiating host receives an initialization error as a result. When the workstation service on a host is stopped, the host discontinues broadcasting a negative name registration response when someone else tries to

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

use the name and sends a name release to a NetBIOS Name Server. The NetBIOS name is said to be released and available for use by another host.

- Name Resolution. When a NetBIOS application wants to communicate with another NetBIOS application, the IP address of the NetBIOS application must be resolved. NetBT performs this function by either broadcasting a NetBIOS name query on the local network or sending a NetBIOS name query to a NetBIOS Name Server.

The NetBIOS name service uses UDP port 137.

### NetBIOS Datagrams

The NetBIOS datagram service provides delivery of datagrams that are connectionless, unsequenced, and unreliable. Datagrams can be directed to a specific NetBIOS name or broadcast to a group of names. Delivery is unreliable in that only the users who are logged on to the network receive the message. The datagram service can initiate and receive both broadcast and directed messages. The NetBIOS datagram service uses UDP port 138.

### NetBIOS Sessions

The NetBIOS session service provides delivery of NetBIOS messages that are connection-oriented, sequenced, and reliable. NetBIOS sessions use TCP connections and provide session establishment, keepalive, and termination. The NetBIOS session service allows concurrent data transfers in both directions using TCP port 139.

### IPv4 Addressing

For IP version 4, each TCP/IP host is identified by a logical IP address. The IP address is a Network layer address and has no dependence on the Data-Link layer address (such as a MAC address of a network adapter). A unique IP address is required for each host and network component that communicates using TCP/IP and can be assigned manually or by using Dynamic Host Configuration Protocol (DHCP).

The IP address identifies a system's location on the network in the same way a street address identifies a house on a city block. Just as a street address must identify a unique residence, an IP address must be globally unique to the internetwork and have a uniform format.

Each IP address includes a network ID and a host ID.

- The network ID (also known as a network address) identifies the systems that are located on the same physical network bounded by IP routers. All systems on the same physical network must have the same network ID. The network ID must be unique to the internetwork.
- The host ID (also known as a host address) identifies a workstation, server, router, or other TCP/IP host within a network. The host address must be unique to the network ID.

### IPv4 Address Syntax

An IP address consists of 32 bits. Instead of expressing IPv4 addresses 32 bits at a time using binary notation (Base2), it is standard practice to segment the 32 bits of an IPv4 address into four 8-bit fields called octets. Each octet is converted to a decimal number (base 10) from 0–255 and separated by a period (a dot). This format is called dotted decimal notation. The following table provides an example of an IP address in binary and dotted decimal formats.

# Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)


## An IP Address in Binary and Dotted Decimal Formats

Binary Format	Dotted Decimal Notation
11000000 10101000 00000011 00011000	192.168.3.24

For example, the IPv4 address of 11000000101010000000001100011000 is:

- Segmented into 8-bit blocks: 11000000 10101000 00000011 00011000.
- Each block is converted to decimal: 192 168 3 24
- The adjacent octets are separated by a period: 192.168.3.24.

The notation w.x.y.z is used when referring to a generalized IP address, and is shown the following figure.



### Types of IPv4 Addresses

The Internet standards define the following types of IPv4 addresses:

- **Unicast.** Assigned to a single network interface located on a specific subnet on the network and used for one-to-one communications.
- **Multicast.** Assigned to one or more network interfaces located on various subnets on the network and used for one-to-many communications.
- **Broadcast.** Assigned to all network interfaces located on a subnet on the network and used for one-to-everyone-on-a-subnet communications.

The following sections describe these types of addresses in detail.

### IPv4 Unicast Addresses

The IPv4 unicast address identifies an interface's location on the network in the same way a street address identifies a house on a city block. Just as a street address must identify a unique residence, an IPv4 unicast address must be globally unique to the network and have a uniform format.

Each IPv4 unicast address includes a network ID and a host ID.

- **The network ID** (also known as a network address) is the fixed portion of an IPv4 unicast address that identifies the set of interfaces that are located on the same physical or logical network segment as bounded by IPv4 routers. A network segment on TCP/IP networks is also known as a subnet. All systems on the same physical or logical subnet must use the same network ID and the network ID must be unique to the entire TCP/IP network.
- **The host ID** (also known as a host address) is the variable portion of an IPv4 unicast address that is used to identify a network node's interface on a subnet. The host ID must be unique to the network ID.

# Windows 2003 TCP/IP Technical Reference

## (Microsoft Corporation)

If the network ID is unique to the TCP/IP network and the host ID is unique to the network ID, then the entire IPv4 unicast address consisting of the network ID and host ID is unique to the entire TCP/IP network.

### IPv4 Multicast Addresses

IPv4 multicast addresses are used for single-packet one-to-many delivery. On an IPv4 multicast-enabled intranet, an IPv4 packet addressed to an IPv4 multicast address is forwarded by routers to the subnets on which there are hosts listening to the traffic sent to the IPv4 multicast address. IPv4 multicast provides an efficient one-to-many delivery service for many types of communication.

IPv4 multicast addresses are defined by the class D Internet address class: 224.0.0.0/4. IPv4 multicast addresses range from 224.0.0.0 through 239.255.255.255. IPv4 multicast addresses for the 224.0.0.0/24 address prefix (224.0.0.0 through 224.0.0.255) are reserved for local subnet multicast traffic.

### IPv4 Broadcast Addresses

IPv4 uses a set of broadcast addresses to provide a one-to-everyone on the subnet delivery service. Packets sent to IPv4 broadcast addresses are processed by all the interfaces on the subnet. The following are the different types of IPv4 broadcast addresses:

- **Network broadcast.** Formed by setting all the host bits to 1 for a classful address prefix. An example of a network broadcast address for the classful network ID 131.107.0.0/16 is 131.107.255.255. Network broadcasts are used to send packets to all interfaces of a classful network. IPv4 routers do not forward network broadcast packets.
- **Subnet broadcast.** Formed by setting all the host bits to 1 for a classless address prefix. An example of a network broadcast address for the classless network ID 131.107.26.0/24 is 131.107.26.255. Subnet broadcasts are used to send packets to all hosts of a classless network. IPv4 routers do not forward subnet broadcast packets. For a classful address prefix, there is no subnet broadcast address, only a network broadcast address. For a classless address prefix, there is no network broadcast address, only a subnet broadcast address.
- **All-subnets-directed broadcast.** Formed by setting all the original classful network ID host bits to 1 for a classless address prefix. A packet addressed to the all-subnets-directed broadcast was defined to reach all hosts on all of the subnets of a subnetted class-based network ID. An example of an all-subnets-directed broadcast address for the subnetted network ID 131.107.26.0/24 is 131.107.255.255. The all-subnets-directed broadcast is the network broadcast address of the original classful network ID. IPv4 routers can forward all-subnets directed broadcast packets, however the use of the all-subnets-directed broadcast address is deprecated in RFC 1812.
- **Limited broadcast.** Formed by setting all 32 bits of the IPv4 address to 1 (255.255.255.255). The limited broadcast address is used for one-to-everyone delivery on the local subnet when the local network ID is unknown. IPv4 nodes typically only use the limited broadcast address during an automated configuration process such as Boot Protocol (BOOTP) or DHCP. For example, with DHCP, a DHCP client must use the limited broadcast address for all traffic sent until the DHCP server acknowledges the use of the offered IPv4 address configuration. IPv4 routers do not forward limited broadcast packets.

### Internet Address Classes

The Internet community originally defined address classes to accommodate different types of addresses and networks of varying sizes. The class of address defined which bits were used for the

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

network ID and which bits were used for the host ID. It also defined the possible number of networks and the number of hosts per network. Of five address classes, class A, B, and C addresses were defined for IPv4 unicast addresses. Class D addresses were defined for IPv4 multicast addresses and class E addresses were defined for experimental uses.


## Class A

Class A network IDs were assigned to networks with a very large number of hosts. The high-order bit in a class A address is always set to zero, which makes the address prefix for all class A networks and addresses 0.0.0.0/1 (or 0.0.0.0, 128.0.0.0). The next seven bits (completing the first octet) are used to enumerate class A network IDs. Therefore, address prefixes for class A network IDs have an 8-bit prefix length (/8 or 255.0.0.0). The remaining 24 bits (the last three octets) are used for the host ID. The address prefix 0.0.0.0/0 (or 0.0.0.0, 0.0.0.0) is a reserved network ID and 127.0.0.0/8 (or 127.0.0.0, 255.0.0.0) is reserved for loopback addresses. Out of a total of 128 possible class A networks, there are 126 networks and 16,777,214 hosts per network.

Note: All-Zeros and All-Ones Host IDs are Reserved. When enumerating host IDs for a given network ID, the two host IDs in which all the bits in the host ID are set to 0 (the all-zeros host ID) and all the bits in the host ID is set to 1 (the all-ones host ID) are reserved and cannot be assigned to network node interfaces. Hence, in the calculation above in which there are 24 bits for class A host IDs, the total number of possible host IDs is 16,777,216 (224). When you subtract the two reserved host IDs, the total number of usable host IDs is 16,777,214.

The following figure illustrates the structure of class A addresses.

### Structure of Class A Addresses




## Class B

Class B network IDs were assigned to medium to large-sized networks. The two high-order bits in a class B address are always set to 10, which makes the address prefix for all class B networks and addresses 128.0.0.0/2 (or 128.0.0.0, 192.0.0.0). The next 14 bits (completing the first two octets) are used to enumerate class B network IDs. Therefore, address prefixes for class B network IDs have a 16-bit prefix length (/16 or 255.255.0.0). The remaining 16 bits (last two octets) are used for the host ID. With 14 bits to express class B network IDs and 16 bits to express host IDs, this allows for 16,384 networks and 65,534 hosts per network.

The following figure illustrates the structure of class B addresses.

### Structure of Class B Addresses




# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## Class C

Class C addresses were assigned to small networks. The three high-order bits in a class C address are always set to 110, which makes the address prefix for all class C networks and addresses 192.0.0.0/3 (or 192.0.0.0, 224.0.0.0). The next 21 bits (completing the first three octets) are used to enumerate class C network IDs. Therefore, address prefixes for class C network IDs have a 24-bit prefix length (/24 or 255.255.255.0). The remaining 8 bits (the last octet) are used for the host ID. With 21 bits to express class C network IDs and 8 bits to express host IDs, this allows for 2,097,152 networks and 254 hosts per network.

The following figure illustrates the structure of class C addresses.



## Class D

Class D addresses are reserved for IPv4 multicast addresses. The four high-order bits in a class D address are always set to 1110, which makes the address prefix for all class D addresses 224.0.0.0/4 (or 224.0.0.0, 240.0.0.0).

## Class E

Class E addresses are reserved for experimental use. The high-order bits in a class E address are set to 1111, which makes the address prefix for all class E addresses 240.0.0.0/4 (or 240.0.0.0, 240.0.0.0)

The following table is a summary of the Internet address classes A, B, and C that can be used for IPv4 unicast addresses.

**Internet Address Class Summary**

Class	Value for <i>w</i>	Network ID Portion	Host ID Portion	Network IDs	Host IDs per Network
A	1-126	<i>w</i>	<i>x.y.z</i>	126	16,777,214
B	128-191	<i>w.x</i>	<i>y.z</i>	16,384	65,534
C	192-223	<i>w.x.y</i>	<i>z</i>	2,097,152	254

## Modern Internet Addresses

The Internet address classes are an obsolete unicast address allocation method that proved to be an inefficient way to assign network IDs and addresses to organizations connected to the Internet. For example, a large organization with a class A network ID can have up to 16,777,214 hosts. However, if the organization only uses 70,000 host IDs, then 16,707,214 potential IPv4 unicast addresses for the Internet are wasted.

On the modern-day Internet, IPv4 address prefixes are handed out to organization's based on the organization's actual need for Internet-accessible IPv4 unicast addresses using a method known as

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

Classless Inter-Domain Routing (CIDR). For example, an organization determines that it needs 2,000 Internet-accessible IPv4 unicast addresses. The Internet Corporation for Assigned Names and Numbers (ICANN) or an Internet service provider (ISP) allocates an IPv4 address prefix in which 21 bits are fixed, leaving 11 bits for host IDs. From the 11 bits for host IDs, the organization can create 2,032 possible IPv4 unicast addresses.

CIDR-based address allocations typically start at 8 bits. The following table lists the required number of host IDs and the corresponding prefix length for CIDR-based address allocations.

**Host IDs Needed and CIDR-based Prefix Lengths**

Number of Host IDs	Prefix Length	Dotted Decimal
2–254	/24	255.255.255.0
255–510	/23	255.255.254.0
511–1,022	/22	255.255.252.0
1,021–2,046	/21	255.255.248.0
2,047–4,094	/20	255.255.240.0
4,095–8,190	/19	255.255.224.0
8,191–16,382	/18	255.255.192.0
16,383–32,766	/17	255.255.128.0
32,767–65,534	/16	255.255.0.0

### Public and Private Addresses

If you want direct (routed) connectivity to the Internet, then you must use public addresses. If you want indirect (proxied or translated) connectivity to the Internet, you can use either public or private addresses. If your intranet is not connected to the Internet in any way, you can use any unicast IPv4 addresses that you want. However, you should use private addresses to avoid network renumbering when your intranet is eventually connected to the Internet.

### Public Addresses

Public addresses are assigned by ICANN and consist of either historically allocated class-based network IDs or, more recently, CIDR-based address prefixes that are guaranteed to be globally unique on the Internet. For CIDR-based address prefixes, the value of w (the first octet) is in the ranges of 1 through 126 and 128 through 223, with the exception of the private address prefixes described in “Private Addresses.”

When the public addresses are assigned, routes are added to the routers of the Internet so that traffic sent to an address that matches the assigned public address prefix can reach the assigned organization. For example, when an organization is assigned an address prefix in the form of a network ID and prefix length, that address prefix also exists as a route in the routers of the Internet. IPv4 packets destined to an address within the assigned address prefix are routed to the proper destination.

### Private Addresses

Each IPv4 interface requires an IPv4 address that is globally unique to the IPv4 network. In the case of the Internet, each IPv4 interface on a subnet connected to the Internet requires an IPv4 address that is globally unique to the Internet. As the Internet grew, organizations connecting to the Internet

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

required a public address for each interface on their intranets. This requirement placed a huge demand on the pool of available public addresses.

When analyzing the addressing needs of organizations, the designers of the Internet noted that for many organizations, most of the hosts on an organization's intranet did not require direct connectivity to the Internet. Those hosts that did require a specific set of Internet services, such as Web access and e-mail, typically access the Internet services through Application layer gateways such as proxy servers and e-mail servers. The result is that most organizations only required a small amount of public addresses for those nodes (such as proxies, servers, routers, firewalls, and translators) that were directly connected to the Internet.

For the hosts within the organization that do not require direct access to the Internet, IPv4 addresses that do not duplicate already-assigned public addresses are required. To solve this addressing problem, the Internet designers reserved a portion of the IPv4 address space and named this space the private address space. An IPv4 address in the private address space is never assigned as a public address. IPv4 addresses within the private address space are known as private addresses. Because the public and private address spaces do not overlap, private addresses never duplicate public addresses.

The private address space specified in RFC 1918 is defined by the following address prefixes:

- **10.0.0.0/8 (10.0.0.0, 255.0.0.0):** Allows the following range of valid IPv4 unicast addresses: 10.0.0.1 to 10.255.255.254. The 10.0.0.0/8 address prefix has 24 host bits that can be used for any addressing scheme within the private organization.
- **172.16.0.0/12 (172.16.0.0, 255.240.0.0):** Allows the following range of valid IPv4 unicast addresses: 172.16.0.1 to 172.31.255.254. The 172.16.0.0/12 address prefix has 20 host bits that can be used for any addressing scheme within the private organization.
- **192.168.0.0/16 (192.168.0.0, 255.255.0.0):** Allows the following range of valid IPv4 unicast addresses: 192.168.0.1 to 192.168.255.254. The 192.168.0.0/16 address prefix has 16 host bits that can be used for any addressing scheme within the private organization.

Because the IPv4 addresses in the private address space will never be assigned by ICANN to an organization connected to the Internet, there will never be routes for the private address prefixes in Internet routers. You cannot connect to a private address over the Internet. Therefore, a host that has a private address must send its Internet traffic requests to an Application layer gateway (such as a proxy server) that has a valid public address or through a network address translator (NAT) that translates the private address into a valid public address.

### Illegal Addresses

Private organization intranets that do not need an Internet connection can choose any address scheme they want, even using public address prefixes that have been assigned by ICANN. If that organization later decides to connect to the Internet, its current address scheme might include addresses already assigned by ICANN to other organizations. These addresses conflict with existing public addresses assigned by ICANN and are known as illegal addresses. Connectivity from illegal addresses to Internet locations is not possible because the routers of the Internet send traffic destined to ICANN-allocated address prefixes to the assigned organizations, not to the organizations using illegal addresses.

For example, a private organization chooses to use the 206.73.118.0/24 address prefix for its intranet. The public address prefix 206.73.118.0/24 has been assigned by ICANN to the Microsoft Corporation

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

and routes exist on the Internet routers to send all packets for IPv4 addresses on 206.73.118.0/24 to Microsoft routers. As long as the private organization does not connect to the Internet, there is no problem because the two address prefixes are on separate IPv4 networks; therefore they are unique to each separate network. If the private organization later connects directly to the Internet and continues to use the 206.73.118.0/24 address prefix, any Internet response traffic to locations matching the 206.73.118.0/24 address prefix is sent to Microsoft routers, not to the routers of the private organization.

### Automatic Private IP Addressing

An interface on a computer running Windows Server 2003 and Windows XP that is configured to obtain an IPv4 address configuration automatically that does not successfully contact a Dynamic Host Configuration Protocol (DHCP) server uses its alternate configuration, as specified on the Alternate Configuration tab.

If the Automatic Private IP Address option is selected on the Alternate Configuration tab and a DHCP server cannot be found, Windows TCP/IP uses Automatic Private IP Addressing (APIPA). Windows TCP/IP randomly selects an IPv4 address from the 169.254.0.0/16 address prefix and assigns the subnet mask of 255.255.0.0. This address prefix has been reserved by the ICANN and is not reachable on the Internet. APIPA allows single-subnet Small Office/Home Office (SOHO) networks to use TCP/IP without static configuration or the administration of a DHCP server. APIPA does not configure a default gateway. Therefore, only local subnet traffic is possible.

### Special IPv4 Addresses

The following are special IPv4 addresses:

- **0.0.0.0:** Known as the unspecified IPv4 address, it is used to indicate the absence of an address. The unspecified address is used only as a source address when the IPv4 node is not configured with an IPv4 address configuration and is attempting to obtain an address through a configuration protocol such as Dynamic Host Configuration Protocol (DHCP).
- **127.0.0.1:** Known as the IPv4 loopback address, it is assigned to an internal loopback interface, enabling a node to send packets to itself.

### Unicast IPv4 Addressing Guidelines

When assigning network IDs to the subnets of an organization, use the following guidelines:

- The network ID must be unique on the IPv4 network. If the network ID is for a subnet on which there are hosts that are directly accessible from the Internet, you must use a public IPv4 address prefix assigned by ICANN or an Internet service provider. If the network ID is for a subnet that is not directly accessible by the Internet, use either a legal public address prefix or a private address prefix that is unique on your private intranet.
- The network ID cannot begin with the numbers 0 or 127. Both of these values for the first octet are reserved and cannot be used for IPv4 unicast addresses.

When assigning host IDs to the interfaces of nodes on an IPv4 subnet, use the following guidelines:

- The host ID must be unique on the subnet.
- You cannot use the all-zeros or all-ones host IDs.

When defining the range of valid IPv4 unicast addresses for a given address prefix, use the following standard practice:

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

- For the first IPv4 unicast address in the range, set all the host bits in the address to 0, except for the low-order bit, which is set to 1.
- For the last IPv4 unicast address in the range, set all the host bits in the address to 1, except for the low-order bit, which is set to 0.

For example, to express the range of addresses for the address prefix 192.168.16.0/20:

- The first IPv4 unicast address in the range is 11000000 10101000 0001000000000001 (host bits are bold), or 192.168.16.1.
- The last IPv4 unicast address in the range is 11000000 10101000 00011111111111110 (host bits are bold), or 192.168.31.254.

### Name Resolution

While IP is designed to work with the 32-bit IP addresses of the source and the destination hosts, computer users are much better at using and remembering names than IP addresses.

When a name is used as an alias for an IP address, a mechanism must exist for assigning that name to the appropriate IP node — to ensure its uniqueness and resolution to its IP address.

In this section, the mechanisms used for assigning and resolving host names (which are used by Windows Sockets applications), and NetBIOS names (which are used by NetBIOS applications) are discussed.

### Host Name Resolution

A host name is an alias assigned to an IP node to identify it as a TCP/IP host. The host name can be up to 255 characters long and can contain alphabetic and numeric characters and the “-” and “.” characters.” Multiple host names can be assigned to the same host. For Windows Server 2003–based computers, the host name does not have to match the Windows Server 2003 computer name.

Windows Sockets applications, such as Microsoft Internet Explorer, can use one of two values to connect to the destination: the IP address or a host name. When the IP address is specified, name resolution is not needed. When a host name is specified, the host name must be resolved to an IP address before IP-based communication with the desired resource can begin.

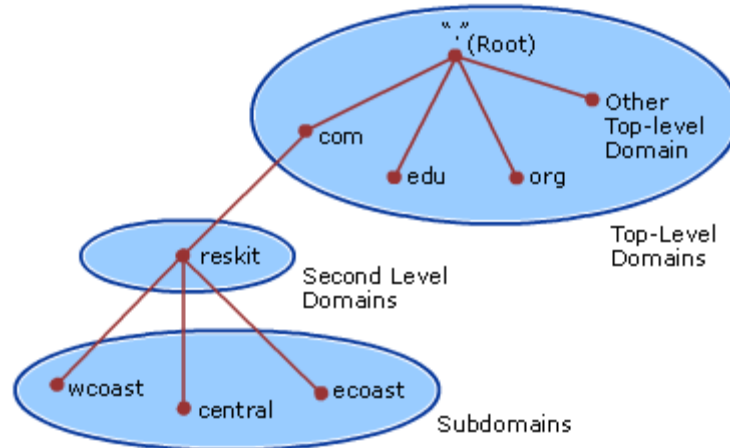
Host names most commonly take the form of a domain name with a structure that follows Internet conventions. Name resolution, and domain names work the same whether they are used for IPv4 or IPv6 addresses.

### Domain Names

To meet the need for a scaleable, customizable naming scheme for a wide variety of organizations, InterNIC has created and maintains a hierarchical namespace called the Domain Name System (DNS). The DNS naming scheme looks like the directory structure for files on a disk. Usually, you trace a file path from the root directory through subdirectories to its final location and its file name. However, a host name is traced from its final location back through its parent domains up to the root. The unique name of the host, representing its position in the hierarchy, is its Fully Qualified Domain Name (FQDN). The top-level domain namespace with second-level and subdomains is shown in the following figure.

# Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

## Domain Name System



The domain namespace includes the following categories:

- The root domain, which is indicated by "" (null), represents the root of the namespace.
- Top-level domains, directly below the root, represent types of organizations. InterNIC is responsible for the maintenance of top-level domain names on the Internet. The following table has a partial list of the Internet's top-level domain names.

### Internet Top-Level Domain Names

Domain Name	Meaning
com	Commercial organization
edu	Educational institution
gov	Government institution
mil	Military group
net	Major network support center
org	Organization other than those above
int	International organization
<country/region code>	Each country/region (geographic scheme)

- Second-level domains, below the top-level domains, represent specific organizations within the top-level domains. InterNIC is responsible for maintaining and ensuring uniqueness of second-level domain names on the Internet.
- Subdomains are below the second-level domain. Individual organizations are responsible for the creation and maintenance of subdomains.

For example, for the FQDN **websrv.wcoast.reskit.com**:

- The trailing period (.) denotes that this is an FQDN with the name relative to the root of the domain namespace. The trailing period is usually not required for FQDNs and if it is missing it is assumed to be present.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

- com is the top-level domain, indicating a commercial organization.
- reskit is the second-level domain, indicating the Resource Kit Corporation.
- wcoast is a subdomain of reskit.com indicating the West Coast division of the Resource Kit Corporation.
- webserv is the name of the Web server in the West Coast division.

Domain names are not case-sensitive.

Organizations not connected to the Internet can implement whatever top and second-level domain names they want. However, typical implementations follow InterNIC specifications so that eventual participation in the Internet will not require a renaming process.

### Host Name Resolution Using a Hosts File

One common way to resolve a host name to an IP address is to use a locally stored database file that contains IP-address-to-host-name mappings. On most UNIX systems, this file is /etc/hosts. On Windows Server 2003 systems, it is the Hosts file in the %systemroot%\System32\Drivers\Etc directory.

The following is an example of the contents of the Hosts file:

```
#  
Table of IP addresses and host names  
#  
127.0.0.1    localhost  
131.107.34.1  router  
172.30.45.121  server1.central.reskit.com s1
```

Within the Hosts file:

- Multiple host names can be assigned to the same IP address. Note that the server at the IP address 172.30.45.121 can be referred to by its FQDN (server1.central.reskit.com) or a nickname (s1). This allows the user at this computer to refer to this server using the nickname s1 instead of typing the entire FQDN.
- Entries can be case sensitive depending on the platform. Entries in the Hosts file for UNIX computers are case-sensitive. Entries in the Hosts file for Windows Server 2003, Windows XP, and Windows 2000–based computers are not case sensitive.

For computers running Windows Server 2003, Windows XP, and Windows 2000, the entries in the Hosts file are loaded into the DNS client resolver cache. When resolving host names, the DNS client resolver cache is always checked.

The advantage of using a Hosts file is that it is customizable for the user. Users can create whatever entries they want, including easy-to-remember nicknames for frequently accessed resources. However, the individual maintenance of the Hosts file does not scale well to storing large numbers of FQDN mappings.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## Host Name Resolution Using a DNS Server

To make host name resolution scalable and centrally manageable, IP address mappings for FQDNs are stored on DNS servers. To enable the querying of a DNS server by a host computer, a component called the DNS resolver is enabled and configured with the IP address of the DNS server. The DNS resolver is a built-in component of TCP/IP protocol stacks supplied with most network operating systems, including Windows Server 2003.

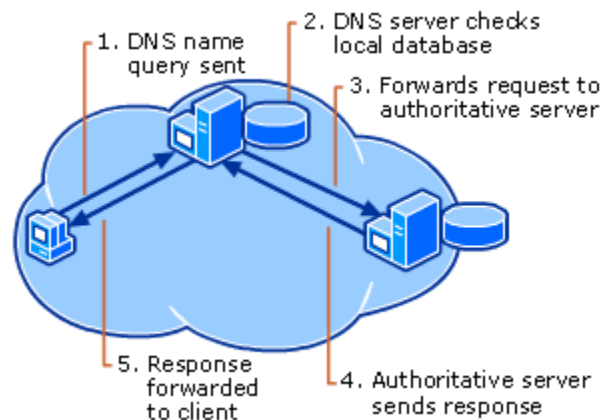
When a Windows Sockets application is given an FQDN as the destination location, the application calls a Windows Sockets function to resolve the name to an IP address. The request is passed to the DNS resolver component in the TCP/IP protocol. The DNS resolver packages the FQDN request as a DNS Name Query packet and sends it to the DNS server.

DNS is a distributed naming system. Instead of storing all the records for the entire namespace on each DNS server, each DNS server stores only the records for a specific portion of the namespace. The DNS server is authoritative for the portion of the namespace that corresponds to records stored on that DNS server. In the case of the Internet, hundreds of DNS servers store various portions of the Internet namespace. To facilitate the resolution of any valid domain name by any DNS server, DNS servers are also configured with pointer records to other DNS servers.

The following process outlines what happens when the DNS resolver component on a host sends a DNS query to a DNS server. This process is shown in the following figure and is simplified so that you can gain a basic understanding of the DNS resolution process.

1. The DNS resolver component of the DNS client formats a DNS Name Query Request message containing the FQDN and sends it to the configured DNS server.
2. The DNS server checks the FQDN in the DNS Name Query Request message against locally stored address records. If a record is found, the IP address corresponding to the requested FQDN is sent back to the client.
3. If the FQDN is not found, the DNS server forwards the request to a DNS server that is authoritative for the FQDN.
4. The authoritative DNS server returns the reply, which contains the resolved IP address, back to the original DNS server.
5. The original DNS server sends the IP address mapping information to the client.

### Resolving an FQDN by using DNS Servers



## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

To obtain the IP address of a server that is authoritative for the FQDN, DNS servers on the Internet go through an iterative process of querying multiple DNS servers until the authoritative server is found. For more information about DNS name-resolution processes, see the DNS Technical Reference.

### Combining a Local Database File with DNS

TCP/IP implementations, including Windows Server 2003, allow the use of both a local database file and a DNS server to resolve host names. When a user specifies a host name in a Windows Sockets–based TCP/IP application: TCP/IP checks the DNS client resolver cache (loaded with entries from the Hosts file and other previously resolved host names) for a matching name. If a matching name is not found in the local database file, the host name is packaged as a DNS Name Query Request message and sent to the configured DNS server.

Combining methods allows the user to have a local database file for resolving personalized nicknames and to use the globally distributed DNS database to resolve FQDNs.

### NetBIOS Name Resolution

NetBIOS name resolution is the process of successfully mapping a NetBIOS name to an IP address. A NetBIOS name is a 16-byte address used to identify a NetBIOS resource on the network. A NetBIOS name is either a unique (exclusive) or group (nonexclusive) name. When a NetBIOS process communicates with a specific process on a specific computer, a unique name is used. When a NetBIOS process communicates with multiple processes on multiple computers, a group name is used.

The NetBIOS name acts as a Session layer application identifier. For example, the NetBIOS session service operates over TCP port 139. All NetBT session requests are addressed to TCP destination port 139. When identifying a NetBIOS application with which to establish a NetBIOS session, the NetBIOS name is used.

An example of a process using a NetBIOS name is the File and Printer Sharing for Microsoft Networks component (the Server service) on a Windows Server 2003–based computer. When you start your computer, the Server service registers a unique NetBIOS name based on your computer's name. The exact name used by the Server service is the 15-character computer name plus a sixteenth character of 0x20. If the computer name is not 15 characters long, it is padded with spaces up to 15 characters long. Other network services, such as the Workstation or Messenger service, also use the computer name to build their NetBIOS names. The sixteenth character is used to uniquely identify each service.

**Note:** The Messenger service referred to here is not Windows Messenger. Windows Messenger is a Microsoft application included in Windows Server 2003 that allows real-time messaging and collaboration.

The Server service on the file server you specify corresponds to a specific NetBIOS name. For example, when you attempt to connect to the computer called CORPSEVER, the NetBIOS name corresponding to the Server service is "CORPSEVER <20>" (note the padding using the space character). Before a file and print sharing connection can be established, a TCP connection must be created. In order for a TCP connection to be established, the NetBIOS name "CORPSEVER <20>" must be resolved to an IP address.

To view the NetBIOS names registered by NetBIOS processes running on a Windows Server 2003 computer, type nbtstat -n at the Windows Server 2003 command prompt.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## NetBIOS Node Types

The exact mechanism by which NetBIOS names are resolved to IP addresses depends on the node's configured NetBIOS Node Type. RFC 1001 defines the NetBIOS Node Types listed in the following table.

### NetBIOS Node Types

Node Type	Description
B-node (broadcast)	B-node uses broadcasted NetBIOS name queries for name registration and resolution. B-node has two major problems: (1) In a large internetwork, broadcasts can increase the network load, and (2) 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 Windows Internet Name Service (WINS), to resolve NetBIOS names. P-node does not use broadcasts; instead, it queries the name server directly. The most significant problem with P-node is that all computers must be configured with the IP address of the NBNS, and if the NBNS is down, computers are not able to communicate even on the local network.
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 it is unable to resolve a name by broadcast, it uses the NBNS of 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 it is unable to resolve a name through the NetBIOS name server, it uses a broadcast to resolve the name.

When NetBT is enabled, Windows Server 2003-based computers are B-node by default and become H-node when configured for a WINS server. Windows Server 2003 also uses a local database file called Lmhosts to resolve remote NetBIOS names.

## IPv4 Routing

After the host name or NetBIOS name is resolved to an IP address, the IP packet must be sent by the sending host to the resolved IP address. Routing is the process of forwarding a packet based on the destination IP address. Routing involves both the TCP/IP host and an IP router. A router is a device that forwards the packets from one network to another. Routers are also commonly referred to as gateways. Both the sending host and router need to make a determination about how the packet is forwarded.

To make these determinations, the IP layer consults a routing table stored in memory. Routing table entries are created by default when TCP/IP initializes and additional entries are added either manually by a system administrator or automatically through communication with routers.

## Direct and Indirect Delivery

IP packets use at least one of two types of delivery based on whether the final destination is located on a directly attached network. These two types of delivery are known as direct and indirect delivery.

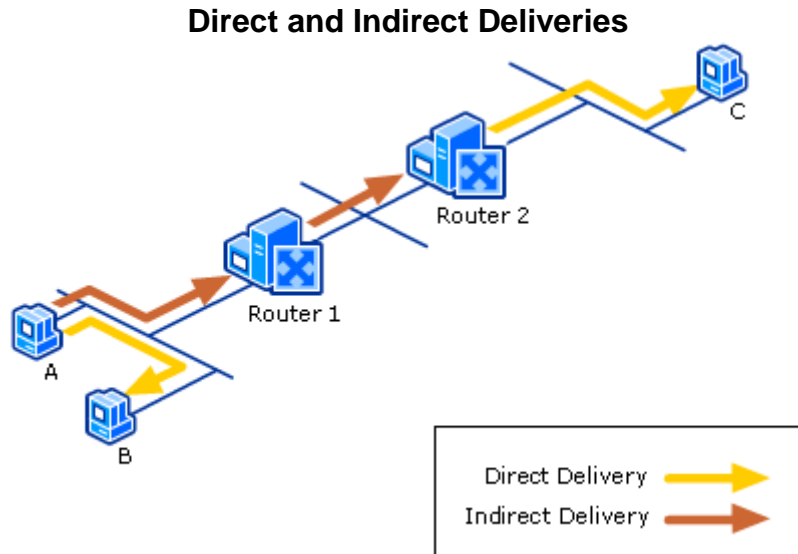
- Direct delivery occurs when the IP node (either the sending node or an IP router) forwards a packet to the final destination on a directly attached network. The IP node encapsulates the IP packet in a frame format for the Network Interface layer (such as Ethernet or Token Ring) addressed to the destination's MAC address.
- Indirect delivery occurs when the IP node (either the sending node or an IP router) forwards a packet to an intermediate node (an IP router) because the final destination is not on a directly

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

attached network. The IP node encapsulates the IP packet in a frame format for the Network Interface layer (such as Ethernet or Token Ring) addressed to the IP router's MAC address.

IP routing is a combination of direct and indirect deliveries.

In the following figure, when sending packets to node B, node A performs a direct delivery. When sending packets to node C, node A performs an indirect delivery to Router 1, and Router 1 performs an indirect delivery to Router 2, and then Router 2 performs a direct delivery to node C.



### IP Routing Table

A routing table is present on all IP nodes. The routing table stores information about IP networks and how they can be reached (either directly or indirectly). Because all IP nodes perform some form of IP routing, routing tables are not exclusive to IP routers. Any node loading the TCP/IP protocol has a routing table. There are a series of default entries according to the configuration of the node and additional entries can be entered either manually through TCP/IP utilities or dynamically through interaction with routers.

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

- **The next-hop IP address.** For a direct delivery, the next-hop IP address is the destination IP address in the IP packet. For an indirect delivery, the next-hop IP address is the IP address of a router.
- **The next-hop interface.** The next-hop interface identifies the physical or logical interface, such as a network adapter, that is used to forward the packet to either its destination or the next router.

### IP Routing Table Entry Types

Entries in the IP routing table contain the following information:

- **Network ID.** The network ID or destination corresponding to the route. The network ID can identify a specific subnet, be a summarized route, or an IP address for a host route. In the Windows Server 2003 IP routing table, this is the Network Destination column.
- **Network mask.** The mask that is used to match a destination IP address to the network ID. In the Windows Server 2003 IP routing table, this is the Netmask column.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

- **Next hop.** The IP address of the next hop. In the Windows Server 2003 IP routing table, this is the Gateway column.
- **Interface.** An indication of which network interface is used to forward the IP packet.
- **Metric.** A number used to indicate the cost of the route so the best route among possible multiple routes to the same destination can be selected. A common use of the metric is to indicate the number of hops (routers crossed) to the network ID.

Entries in the routing table can be used to store the following types of routes:

- **Directly attached network ID.** A route for network IDs that are directly attached. For directly attached networks, the Next Hop field can be blank or contain the IP address of the interface on that network.
- **Remote network ID.** A route for network IDs that are not directly attached but are available across other routers. For remote networks, the Next Hop field is the IP address of a local router.
- **Host route.** 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 the IP address of the specified host and the network mask is 255.255.255.255.
- **Default route.** The default route is designed to be used when a more specific network ID or host route is not found. The default route network ID is 0.0.0.0 with a network mask of 0.0.0.0.

### Route Determination Process

To determine which routing table entry is used to find the next-hop address and interface, IP uses the following process:

- For each entry in a routing table, IP performs a bit-wise logical AND operation between the destination IP address and the network mask. It compares the result with the network ID of the entry for a match.
- A list of matching routes is compiled. The route that has the longest match (the route with the largest number of bits that match the destination IP address) is chosen. The longest matching route is the most direct route to the destination IP address. If multiple matching entries are found (for example, multiple routes to the same network ID), the router uses the lowest metric to select the best route. If multiple entries have the longest match and the lowest metric, the router designates one of them as the routing table entry. For Windows Server 2003 TCP/IP, the route chosen corresponds to the route associated with the interface that is first in the network binding order.

The end result of the route-determination process is a single route in the routing table that yields a next-hop IP address and interface. If the route-determination process fails to find a route, IP indicates a routing error. For the sending host, an IP routing error message is sent to the upper layer protocol, such as TCP or UDP. For a router, an ICMP Destination Unreachable–Host Unreachable message is sent to the sending host.

### Routing Table for Windows Server 2003

The following table shows the default routing table for a Windows Server 2003–based host (not a router). The host has a single network adapter and has the IP address 157.60.27.90, subnet mask 255.255.240.0, and a default gateway of 157.60.16.1.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## Windows Server 2003 Routing Table

Network Destination	Netmask	Gateway	Interface	Metric	Purpose
0.0.0.0	0.0.0.0	157.60.16.1	157.60.27.90	1	Default Route
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1	1	Loopback Network
157.60.16.0	255.255.240.0	157.60.27.90	157.60.27.90	1	Directly Attached Network
157.60.27.90	255.255.255.255	127.0.0.1	127.0.0.1	1	Local Host
157.60.255.255	255.255.255.255	157.60.27.90	157.60.27.90	1	Network Broadcast
224.0.0.0	240.0.0.0	157.60.27.90	157.60.27.90	1	Multicast
255.255.255.255	255.255.255.255	157.60.27.90	157.60.27.90	1	Limited Broadcast

### Default Route

The entry corresponding to the default gateway configuration is a network destination of 0.0.0.0 with a network mask (netmask) of 0.0.0.0. Any destination IP address that is logically ANDed with 0.0.0.0 results in 0.0.0.0. Therefore, for any IP address, the default route produces a match. If the default route is chosen because no better routes were found, the IP packet is forwarded to the IP address in the Gateway column (the default gateway of 157.60.16.1), by using the interface assigned the IP address in the Interface column.

### Loopback Network

The loopback network entry is designed to take any IP address of the form 127.x.y.z and forward it to the special loopback address of 127.0.0.1.

### Directly Attached Network

The local network entry corresponds to the directly attached network. IP packets destined for the directly attached network are not forwarded to a router but sent directly to the destination. Note that the Gateway and Interface columns match the IP address of the node. This indicates that the packet is sent from the network adapter corresponding to the node's IP address.

### Local Host

The local host entry is a host route (network mask of 255.255.255.255) corresponding to the IP address of the host. All IP packets sent to the IP address of the host are forwarded to the loopback address.

### Network Broadcast

The network broadcast entry is a host route (network mask of 255.255.255.255) corresponding to the all-subnets directed broadcast address (all subnets of class B network ID 157.60.0.0). Packets addressed to the all-subnets directed broadcast are sent from the network adapter corresponding to the node's IP address.

### Multicast

The multicast addresses route is used to send any multicast IP packets from the network adapter corresponding to the node's IP address.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## Limited Broadcast

The limited broadcast address is a host route (network mask of 255.255.255.255). Packets addressed to the limited broadcast are sent from the network adapter corresponding to the node's IP address.

Viewing the IP Routing Table

To view the IP routing table on a Windows Server 2003-based computer, type `route print` at a Windows Server 2003 command prompt.

When determining the next-hop IP address and interface from a route in the routing table:

- If the gateway address is the same as the interface address, the next-hop IP address is set to the destination IP address of the IP packet.
- If the gateway address is not the same as the interface address, the next-hop IP address is set to the gateway address.

For example, when traffic is sent to 157.60.16.48, the most specific matching route is the route for the directly attached network (157.60.16.0/20). The next-hop IP address is set to the destination IP address (157.60.16.48) and the interface is the network adapter that has been assigned the IP address 157.60.27.90.

When sending traffic to 192.168.0.79, the most specific matching route is the default route (0.0.0.0/0). The next-hop IP address is set to the gateway address (157.60.16.1) and the interface is the network adapter that has been assigned the IP address 157.60.27.90.

## Maintenance of Routing Table Entries

For IP routing to occur efficiently between routers in the IP internetwork, routers must be configured with remote network IDs or a default route. On large IP internetworks, one of the challenges faced by network administrators is how to maintain the routing tables on their IP routers so that IP traffic flow is traveling the best path and is fault tolerant.

There are two methods of maintaining routing table entries on IP routers.

### Manual

Static IP routers have routing tables that do not change unless manually changed by a network administrator.

Static routing relies on the manual administration of the routing table. Remote network IDs are not discovered by static routers and must be manually configured. Static routers are not fault-tolerant. If a static router goes down, neighboring routers do not sense the fault and inform other routers.

### Automatic

Dynamic IP routers have routing tables that change automatically based on the exchange of routing information with other routers.

Dynamic routing employs the use of routing protocols, such as Routing Information Protocol (RIP) and Open Shortest Path First (OSPF), to dynamically update the routing table through the exchange of routing information between routers. Remote network IDs are discovered by dynamic routers and automatically entered into the routing table. Dynamic routers are fault-tolerant. If a dynamic router goes down, the fault is detected by neighboring routers, which send the changed routing information to the other routers in the internetwork.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## Physical Address Resolution

Based on the destination IP address and the route determination process, IP determines the next-hop IP address and interface. IP then sends the IP packet, the next-hop IP address, and the interface to ARP.

If the next-hop IP address is the same as the destination IP address, then ARP performs a direct delivery. In a direct delivery, the MAC address corresponding to the destination IP address must be resolved.

If the next-hop IP address is not the same as the destination IP address, then ARP performs an indirect delivery. The next-hop IP address is the IP address of a router between the current IP node and the final destination. In an indirect delivery, the MAC address corresponding to the IP address of the router must be resolved.

To resolve a next-hop IP address to its MAC address, ARP uses broadcast traffic on shared access networking media (such as Ethernet or Token Ring) to send out a broadcasted ARP Request frame. An ARP Reply, containing the MAC address corresponding to the requested next-hop IP address, is sent back to the sender of the ARP Request.

## ARP Cache

To keep the number of broadcasted ARP Request frames to a minimum, many TCP/IP protocol stacks incorporate an ARP cache, which is a table of recently resolved IP addresses and their corresponding MAC addresses. TCP/IP checks the ARP cache before sending an ARP Request frame. Each interface has its own ARP cache.

Depending on the vendor implementation, the ARP cache can have the following qualities:

- ARP cache entries can be dynamic (based on ARP Replies) or static. Static ARP entries are permanent and are manually added by using a TCP/IP utility such as the ARP tool provided with Windows Server 2003. Static ARP cache entries are used to prevent ARP Requests for commonly used local IP addresses, such as routers and servers. The problem with static ARP entries is that they have to be manually updated when network interface equipment changes.
- Dynamic ARP cache entries have a time-out value associated with them to remove entries in the cache after a specified period of time. Dynamic ARP cache entries for Windows Server 2003 TCP/IP are given a maximum time of 10 minutes before being removed.

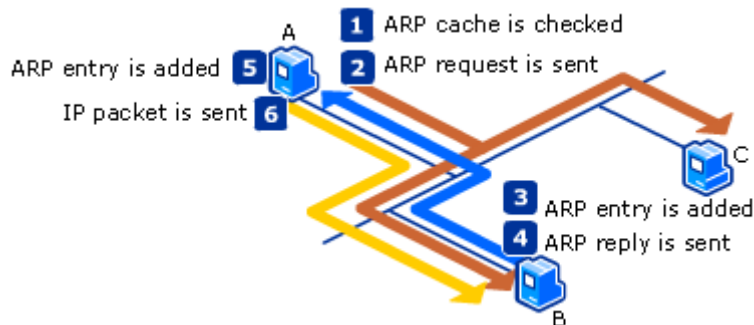
To view the ARP cache on a Windows Server 2003-based computer, type `arp -a` at a Windows Server 2003 command prompt.

## ARP Process

IP sends ARP the IP packet, the next-hop IP address, and the next-hop interface. Whether performing a direct or indirect delivery, ARP carries out the following process, as shown in the following figure.

# Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

## ARP Process



1. Based on the next-hop address and interface, ARP consults the appropriate ARP cache for an entry for the next-hop IP address. If an entry is found, ARP skips to step 6.
2. If an entry is not found, ARP builds an ARP Request frame containing the MAC address of the interface sending the ARP Request, the IP address of the interface sending the ARP Request, and the next-hop IP address. ARP then broadcasts the ARP Request using the appropriate interface.
3. All hosts receive the broadcasted frame and the ARP Request is processed. If the receiving host's IP address matches the requested IP address (the next-hop IP address), its ARP cache is updated with the address mapping of the sender of the ARP Request.
4. If the receiving host's IP address does not match the requested IP address, the ARP Request is silently discarded.
5. The receiving host formulates an ARP Reply containing the requested MAC address and sends it directly to the sender of the ARP Request.
6. When the ARP Reply is received by the sender of the ARP Request, it updates its ARP cache with the address mapping.
7. The ARP Request host and the ARP Reply host have each other's address mappings in their ARP caches.
8. ARP sends the IP packet to the next-hop node by addressing it to the resolved MAC address.

### End-to-End Delivery

The IP routing processes for all nodes involved in the delivery of an IP packet include the sending host, the intermediate routers, and the destination host.

### IP on the Sending Host

When a host sends a packet, the packet is transmitted from an upper layer protocol (TCP, UDP, or ICMP) to IP, and then IP on the sending host does the following:

1. Sets the Time-to-Live (TTL) value to either a default or application-specified value.
2. Checks its routing table for the best route to the destination IP address. If no route is found, IP sends a routing error message to the upper layer protocol (TCP, UDP, or ICMP).
3. Determines the next-hop IP address and the interface, based on the most specific matching route.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

4. Sends the packet, the next-hop IP address, and the next-hop interface to Address Resolution Protocol (ARP), and then ARP resolves the next-hop IP address to its media access control (MAC) address and forwards the packet.

### IP on the Router

When a packet is received at a router, the packet is passed to IP, and IP on the router does the following:

1. Verifies the IP header checksum. If the IP header checksum fails, the IP packet is discarded without notification to the user. This is known as a silent discard.
2. Verifies whether the destination IP address in the IP packet corresponds to an IP address assigned to a router interface. If so, the router processes the IP packet as the destination host (see step 3 in the following "IP on the Destination Host" section).
3. If the destination IP address is not the router, IP decrements the Time-to-Live (TTL). If the TTL is 0, the router discards the packet and sends an ICMP Time Expired–TTL Expired in Transit message to the sender.
4. If the TTL is 1 or greater, IP updates the TTL field and calculates a new IP header checksum.
5. IP checks its routing table for the best route to the destination IP address in the IP packet. If no route is found, the router discards the packet and sends an ICMP Destination Unreachable–Host Unreachable message to the sender.
6. Based on the best route found, IP determines the next-hop IP address and interface.
7. IP sends the packet, the next-hop IP address, and the interface to ARP, and then ARP forwards the packet to the appropriate MAC address.

This entire process is repeated at each router in the path between the source and destination host.

### IP on the Destination Host

When a packet is received at the destination host, it is passed up to IP, and IP on the destination host does the following:

1. Verifies the IP header checksum. If the IP header checksum fails, the IP packet is silently discarded.
2. Verifies that the destination IP address in the IP packet corresponds to an IP address assigned to the host. If the destination IP address is not assigned to the host, the IP packet is silently discarded.
3. Passes the IP packet without the IP header to the appropriate upper-level protocol, based on the IP protocol field. If the protocol does not exist, ICMP sends a Destination Unreachable–Protocol Unreachable message back to the sender.
4. For TCP and UDP packets, IP checks the destination port and processes the TCP segment or UDP header. If no application exists for the UDP port number, ICMP sends a Destination Unreachable–Port Unreachable message back to the sender. If no application exists for the TCP port number, TCP sends a Connection Reset segment back to the sender.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## TCP/IP Tools and Settings

Microsoft Windows Server 2003 and Microsoft Windows XP Professional support many TCP/IP diagnostic and remote networking tools that help manage and troubleshoot your network. Windows Server 2003 adds new IPv6 tools and updates several standard IPv4 tools to include IPv6 functionality.

## TCP/IP Tools

The following sections “TCP/IP Diagnostic Tools” and “TCP/IP Remote Networking Tools” describe the diagnostic and remote networking tools associated with TCP/IP.

## Other Tools

In addition to the Microsoft TCP/IP tools described in this document, Windows Server 2003 supports the tools shown in the following table. The Ipsec6 and Ttcp tools are new IPv6 tools.

### Other TCP/IP-Related Tools

Tool	Description
Ipsec6	You can use <b>Ipsec6</b> to experiment with managing Internet Protocol security (IPSec) in an IPv6 test environment by configuring IPSec policies and security associations.
Ttcp	You can use this tool to troubleshoot TCP and UDP traffic. You can find the <b>Ttcp</b> tool in the <b>Valueadd\Msft\Net\Tools</b> folder of the Windows Server 2003 CD-ROM.
Event Viewer	Tracks errors and events recorded in Application, Security, and System logs.
Network Monitor	Captures and displays network traffic. The full version, which can capture all frames sent on a network segment, is part of the Microsoft Systems Management Server (SMS) product. A limited version, which can capture frames that are sent to or from the network adapter of the computer on which Network Monitor is installed, is included with Windows Server 2003.
Registry Editor	The registry editor Regedit.exe allows viewing and editing of registry parameters.
Simple Network Management Protocol (SNMP) service	Provides statistical information about network devices to SNMP management systems, such as HP OpenView, IBM NetView, or Sun Net Manager.
System Monitor	Analyzes TCP/IP network performance. System Monitor is a component of in the Performance Logs and Alerts snap-in.

## TCP/IP Diagnostic Tools

TCP/IP supports several diagnostic tools in Windows Server 2003 that include new support for IPv6:

- Netstat
- Pathping
- Ping
- Tracert

In addition, Windows Server 2003 introduces the Netsh commands for both IPv6 and for Ipv6 6to4. You can use the first set of tools to query and configure IPv6 interfaces, addresses, caches, and routes. You can use the second set of commands to query or configure the 6to4 service on either a 6to4 host or a 6to4 router.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

The following table lists TCP/IP diagnostic tools that you can use to identify and resolve TCP/IP networking problems.

### TCP/IP Diagnostic Tools

Tool	Description
Arp	View and manage the Address Resolution Protocol (ARP) cache on the interfaces of the local computer.
Hostname	Display the name of the computer.
Ipconfig	Display current TCP/IP network configuration values, update or release Dynamic Host Configuration Protocol (DHCP) allocated leases, and display, register, or flush Domain Name System (DNS) names.
Nbtstat	Check the state of current NetBIOS over TCP/IP (NetBT) connections, view and update the NetBIOS name cache, and determine the names registered with Windows Internet Name Service (WINS).
Netdiag	Check the state of a network client to help isolate TCP/IP-related connectivity problems, including verifying that DNS is available and functioning correctly.
Netsh	Provides thirteen sets of commands (called contexts) for performing a wide range of network configuration tasks. Windows Server 2003 adds a new context for managing IPv6 to the <b>netsh</b> command set.
Netstat	Display statistics for current TCP/IP connections. Windows Server 2003 adds IPv6 parameters to the <b>netstat</b> command.
Nslookup	Check records, domain host aliases, domain host services, and operating system information by querying DNS servers.
Pathping	Trace a path to a remote system and report packet losses at each router along the way. Windows Server 2003 adds IPv6 parameters to the <b>pathping</b> command.
Ping	Send Internet Control Message Protocol (ICMP) Echo messages to verify IP connectivity. Windows Server 2003 adds IPv6 parameters to the <b>ping</b> command.
Route	Display the IP routing table, and add, edit, or delete IPv4 routes. Route for Windows Server 2003 also displays IPv6 routes.
Tracert	Trace a path to a destination. Windows Server 2003 adds IPv6 parameters to the <b>tracert</b> command.

All of the tools in the previous table, except Netdiag, are installed with the Windows Server 2003 operating system. Netdiag is one of the support tools delivered with the second Windows Server 2003 operating system CD, which you install separately from the operating system itself.

#### Arp.exe: Arp

The Arp command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Arp command-line tool to display and modify entries in the local Address Resolution Protocol (ARP) cache. The ARP cache, which is a memory-resident list, contains one or more tables that store IP addresses and the corresponding Ethernet, Token Ring, or wireless LAN physical addresses that have been resolved from other computers on the same subnet. Typically, a physical address is the Media Access Control (MAC) address. A separate table exists for each network adapter installed on the computer.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## **Hostname.exe: Hostname**

The Hostname command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Hostname command-line tool to display the host name of the computer on which you run the command. Hostname does not display the fully qualified domain name (FQDN) of the computer. You can find the computer name and its FQDN through Control Panel on the computer.

## **Ipconfig.exe: Ipconfig**

The Ipconfig command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Ipconfig command-line tool to display the current configuration of the installed IP stack on a networked computer and to refresh DHCP and DNS settings. The Ipconfig command is often one of the first commands you use to check the status of the connection when you experience communication problems on a TCP/IP network. Ipconfig is most useful for managing computers that obtain an IP address automatically, such as by using DHCP or through alternate configuration.

When used without a parameter, Ipconfig displays the IPv4 address, subnet mask, and default gateway for all adapters on a computer. If the computer has the IPv6 protocol installed, Ipconfig also displays the IPv6 address information.

## **Nbtstat.exe: Nbtstat**

The Nbtstat command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Nbtstat command-line tool to troubleshoot NetBIOS name-resolution problems.

When a network is functioning correctly, NetBIOS over TCP/IP (NetBT) resolves NetBIOS names to IP addresses. NetBT uses several options for NetBIOS name resolution, including local NetBIOS name cache lookup, WINS server query, broadcast, and LMHOSTS lookup. After NetBT methods are exhausted, Windows TCP/IP converts the NetBIOS name to a host name and attempts host name resolution, including checking the local host name, checking the DNS client resolver cache, and querying DNS servers.

Use Nbtstat to display a variety of information, including:

- NetBT protocol statistics.
- NetBIOS name tables for both the local computer and for remote computers. The NetBIOS name table is the list of NetBIOS names that corresponds to NetBIOS applications running on that computer.
- NetBIOS name cache. The NetBIOS name cache is the table that contains NetBIOS name-to-IP address mappings.
- Also use Nbtstat to refresh the NetBIOS name cache and the names registered with WINS.

## **Netdiag.exe: Netdiag**

Netdiag is a command-line tool delivered with the support tools on the second Windows Server 2003 operating system CD, which you install separately from the operating system itself. This tool is included with all versions of Windows that include TCP/IP.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

Use the Netdiag command-line tool, which is one of the support tools delivered with the second Windows Server 2003 operating system CD, to help isolate TCP/IP-related connectivity problems. After you install the Support Tools, Netdiag.exe appears in the Program Files\Support Tools folder on the drive on which the Support Tools are installed.

Netdiag provides tests that report information about a computer and its network configuration, which can help diagnose network problems. Use the Netdiag logging parameter (-l) to capture and store output from Netdiag tests. One use of Netdiag logging is to add the output of Netdiag tests to your network's baseline documentation and to create updated test logs each time important changes to a computer's configuration are made.

In addition to troubleshooting TCP/IP issues, you can also use Netdiag to examine a computer's Internetwork Packet Exchange (IPX) and NetWare configurations.

**Note:** The IPX protocol is supported only on the 32-bit platforms in Windows Server 2003.

Despite the name similarity and some overlap in function, the Netdiag command-line tool that you install from the Windows Server 2003 operating system CD as part of the set of Support Tools is not the same as the following two versions of a different network diagnostic tool:

Network Diagnostics, which you access under Help and Support Center Tools in Help and Support Center for Windows Server 2003.

Netsh diagnostic (diag)commands, which you access by typing netsh -c diag at a command prompt.

Network Diagnostics and the Netshdiag commands use the same Dynamic Link Library (DLL) and run the same set of tests. However, Netsh lets you test user-defined destinations, whereas Network Diagnostics does not.

Although there is some overlap between these tools and the Netdiag support tool described here, for the most part Netdiag provides a different set of tests.

### **Netsh.exe: Netsh**

The Netsh command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Netsh command-line tool to locally or remotely display or modify the configuration of services or protocols on computers running Windows. In addition, the Netsh command-line interface is scriptable, which lets you perform batch configurations or administration from a centralized location.

The Netsh commands, which were first introduced in Windows 2000, help you manage a standard TCP/IP network, that is, an IPv4 network. For organizations that have begun to explore IPv6 by introducing IPv6 into their network, Windows Server 2003 and Windows XP add a new set of Netsh commands for IPv6.

For administrative convenience, the Netsh commands are grouped into sets, called contexts. Each context provides commands appropriate for a specific area of networking functionality. A context is implemented through an associated Netsh helper, which is a dynamic link library (DLL) file that provides the capabilities for that context. Netsh directs the context command that you enter to the appropriate helper, and the helper then carries out the command. The Netsh helper DLLs interact with other operating system components, such as WINS or TCP/IP. For example, Winsmon.dll provides

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

the set of commands for managing the WINS service, and Ifmon.dll provides the set of commands for managing IPv4 interfaces.

A context can contain one or more additional contexts, called a subcontext. For example, the IPv6 context contains the 6to4 subcontext, which provides a separate set of commands for managing the 6to4 service (which encapsulates IPv6 traffic with an IPv4 header before it is sent over an IPv4 network) on either a 6to4 host or a 6to4 router.

Netsh is extensible. Developers can create additional contexts to manage networking services in addition to the contexts provided by Windows.

### **Netstat.exe: Netstat**

The Netstat command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Netstat command-line tool to display active TCP/IP connections, ports on which the computer is listening, Ethernet statistics, the IP routing table, IPv4 statistics (for the IP, ICMP, TCP, and UDP protocols), and IPv6 statistics (for the IPv6, ICMPv6, TCP over IPv6, and UDP over IPv6 protocols).

### **Nslookup.exe: Nslookup**

The Nslookup command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Nslookup command-line tool for querying and troubleshooting the DNS infrastructure. For example, Nslookup can provide host-name resolution.

### **Pathping.exe: Pathping**

The Pathping command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use Pathping, a command-line tool that combines features of Ping and Tracert, to obtain additional information that neither of those tools provides. Specifically, you can use Pathping to identify the route to a remote host, then ping the remote host for a period of time to collect and report statistics. Pathping information includes information about the intermediate routers visited on the path, the Round-Trip Time (RTT) value, and link-loss information.

### **Ping.exe: Ping**

The Ping command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Ping command-line tool as your primary tool for troubleshooting IP-level connectivity between two TCP/IP computers. Ping sends ICMP Echo or ICMPv6 Echo Request messages to perform network diagnostics and to test ability to reach a specific destination. Ping can use IPv4 or IPv6 addresses. If a name is specified, Ping uses the address that is resolved.

Ping lets you specify the size of packets to use (the default is 32 bytes), how many to send, whether to record the route used, what Time-To-Live (TTL) value to use, whether or not to set the Don't Fragment flag, and so on.

### **Route.exe: Route**

The Route command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

Use the Route command-line tool to view and modify the local IP routing table.

For two hosts to exchange IP datagrams, they must both have a route to each other, or they must use a default gateway that knows a route between the two. Typically, routers exchange information using a protocol such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF). The RIP Listening service is available for Microsoft Windows XP Professional, and full routing protocols are supported by Windows Server 2003 in the Routing and Remote Access service.

All symbolic names used for the specified network destination of the route are looked up in the network database file NETWORKS. The symbolic names for the gateway are looked up in the host name database file HOSTS. If the command is print or delete, the destination value can be a wildcard value specified by an asterisk ("\*"). If the destination specified contains a \* or ?, it is treated as a shell pattern and only matching destination routes are printed. The asterisk matches any string, and the question mark matches any one character. For example, 157.\*.1, 157.\*, 127.\*, and \*224\* are all valid uses of the wildcard asterisk.

Using an invalid combination of a destination and netmask value generates a "The route addition failed: The specified mask parameter is invalid. (Destination & Mask) != Destination." error. This sort of error message appears, for example, when a bitwise logical AND between the destination and mask does not equal the destination value.

### **Tracert.exe: Tracert**

The Tracert command-line tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Tracert command-line route-tracing tool to display the path between the sending host and a destination.

The path that Tracert displays is a list of near-side router interfaces of the routers in the path between the source host and destination. Tracert uses the IP Time-to-Live (TTL) field in Internet Control Message Protocol (ICMP) Echo Requests and ICMP Time Exceeded-TTL Exceeded in Transit messages to determine the path from a source to a destination through an IP internetwork.

Note that some routers silently drop packets with expired TTLs. These routers do not appear in the Tracert display.

Tracert works by incrementing the TTL value by one for each ICMP Echo Request it sends, and then waiting for an ICMP Time Exceeded-TTL Exceeded in Transit message. The TTL values of the Tracert packets start with an initial value of one; the TTL of each trace after the first is incremented by one. A packet sent out by Tracert travels one hop further on each successive trip.

**Note:** The UNIX version of Tracert performs the same function as the Windows version, except that the IP payload is a UDP packet addressed to a (presumably) unknown destination UDP port. Intermediate routers send back ICMP Time Expired-TTL Exceeded in Transit messages recording the route taken, and the final destination sends back an ICMP Destination Unreachable-Port Unreachable message. The UDP payload from the UNIX Tracert tool can cross routers and some firewalls, whereas the ICMP Echo Request messages might not, due to ICMP filtering. To avoid this problem in Windows Server 2003, turn off packet filtering and then try using Tracert again.

# Windows 2003 TCP/IP Technical Reference

(Microsoft Corporation)

## TCP/IP Remote Networking Tools

Microsoft TCP/IP includes several remote networking tools. The following table lists the tools included with Microsoft TCP/IP that you can use to communicate with remote computers.

### TCP/IP Remote Tools

Tool	Description
Finger	Displays information about a user or users on a specified remote computer (typically, a computer running UNIX) that is running the Finger service.
Ftp	Transfers files over the Internet to, and from, a computer running a File Transfer Protocol (FTP) server service, such as the FTP component of Microsoft Internet Information Services (IIS).
Rcp	Copies files between a computer running Windows Server 2003 and a computer running Rshd, the UNIX remote shell service.
Rexec	Runs commands on a remote computer (typically, a computer running UNIX).
Rsh	Runs commands on remote hosts using the Rsh service, the UNIX remote shell service.
Telnet	Starts terminal emulation with a remote host running a Telnet server service.
Tftp	Transfers files to and from a remote computer (typically, a computer running UNIX) that is running the Trivial File Transfer Protocol (TFTP) service.

Note: All passwords used by Windows networking services are encrypted. However, the Ftp, Rexec, and Telnet connectivity tools rely on plaintext password authentication by the remote computer. Plaintext passwords are not encrypted before being sent over the network. This enables another user equipped with a network capture tool such as Network Monitor on the same network to obtain a user's remote account password. For this reason, choose different passwords from those used for computers running Windows Server 2003 or domains when connecting to non-Microsoft remote computers with the Ftp, Rexec, or Telnet tools. Note that the protocols themselves prohibit encryption; the use of plaintext passwords is not recommended by Microsoft.

#### Finger.exe: Finger

The Finger connectivity tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Finger connectivity tool to display information about a user or users on a specified remote computer (typically, a host computer running UNIX) that is running the Finger service. Output varies based on how the remote host specifies Finger output format.

#### Ftp.exe: Ftp

The Ftp connectivity tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

FTP is a protocol that defines how to transfer files from one computer to another over a TCP/IP network, such as the Internet or a company intranet. Use the Ftp connectivity tool to transfer files to and from a host running an FTP server service, such as the FTP component of Microsoft Internet Information Services (IIS). You can use Ftp interactively or in batch mode to process ASCII text files.

**Note:** The FTP service is a component of Microsoft Internet Information Services (IIS). However, when you install IIS on a server, FTP is not installed unless you explicitly specify that it be installed. If

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

you install IIS without FTP, you can use Add or Remove Windows Components in Add or Remove Programs in Control Panel to install FTP later.

Typically, users use Ftp to download files from a location on the Internet or to upload files to a location on the Internet. Before a user can log on to an FTP server, the server administrator must grant the user permission to access the inetput\ftproot folder, which is created along with the Default FTP site by default when you install IIS FTP, or a folder for the use of a particular user or group under the ftproot folder. Alternatively, you can create a new FTP site, and you can use folders other than ftproot and its children for FTP files.

Similarly, users can upload a file from the local computer to an FTP server using the put command. However, for users to successfully use the put command, the administrator must grant them write permissions for the FTP folder to which the file will be uploaded.

### **Rcp.exe: Rcp**

The Rcp connectivity tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Rcp connectivity tool to copy files between a computer running Windows Server 2003 and a computer running Rshd, the UNIX remote shell server service or daemon. Rshd stands for Rsh daemon; daemon is the UNIX term for service. In addition to running Rsh, the remote computer must also support the Rcp tool. Also use Rcp for third-party transfer to copy files between two computers running Rsh when the command is issued from computers running Windows Server 2003. The UNIX Rsh server service is not available on computers running Windows Server 2003, but the computer running Windows Server 2003 can participate as the computer from which the commands are issued.

### **Permitting Network Access**

The .rhosts file typically permits network access on UNIX systems. The .rhosts file lists computer names and associated logon names that have access to remote computers. When you run Rcp, Rexec, or Rsh tools remotely with a correctly configured .rhosts file, you do not need to provide logon and password information for the remote computer. The .rhosts file must be in the user's home directory on the remote computer.

The .rhosts file specifies which remote computers or users can access a local account using the Rcp or Rsh commands (Rsh is described later in this chapter). For you to access the remote computer using Rcp, the .rhosts file (or a file called Hosts.equiv) must exist on the remote computer. Rcp transmits the local user name to the remote computer. The remote computer uses this name and the IP address (usually resolved to a computer name) of the requesting computer to determine whether to grant access. No provision exists for specifying a password to access an account using Rcp. If the user is logged on to a Windows Server 2003 domain, the domain controller must be available to resolve the currently logged-on name because the logged-on name is not cached on the local computer. Because the user name is required as part of the Rcp protocol, the command fails if it cannot obtain the user name.

The .rhosts file is a text file in which each line is an entry. An entry consists of the local host name, the local user name, and any comments about the entry. Each entry is separated by a tab or space, and comments begin with the number sign (#).

### **Specifying Computers (Hosts)**

Use Host.user to specify a user name other than the current user name. If you use Host.user with Source (which specifies which files to copy), the .rhosts file on the remote computer must contain an entry for user. For example:

# Windows 2003 TCP/IP Technical Reference

## (Microsoft Corporation)

```
Rcp host99.user7:file1 corp7.admin:file2
```

In this example, the .rhosts file on Corp7 must have an entry for User7 on Host99.

If you type a computer name as a fully qualified domain name (FQDN) containing periods, you must append a user name to the host name. This prevents the last element of the domain name from being interpreted as a user name. For example:

```
Rcp domain-name1.user:johns domain-name2.user:buddyg
```

### Performing Remote Processing

Remote processing is performed by a command run from the user's logon shell on most UNIX computers. The user's .profile or .cshrc file is run before file names are parsed. You can specify that exported shell variables be used in remote file names by using one of the escape characters. The escape characters are the backslash (\), quotation mark ("), and apostrophe (').

### Copying Files

If you try to copy several files to a file rather than to a directory, only the last file is copied. Also, the rcp command cannot copy a file onto itself (Source and Path/Destination cannot be the same).

### Rexec.exe: Rexec

The Rexec connectivity tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Rexec connectivity tool to run commands on a remote computer that is running the Rexecd service, the UNIX remote execute service. Rexecd stands for Rexec daemon; daemon is the UNIX term for service. Before carrying out the specified command, Rexec authenticates the user name on the remote host by prompting for a password. Windows Server 2003 does not provide the Rexec service.

Rexec copies standard input to the remote command, standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. Interrupt, Quit, and Terminate signals are propagated to the remote command. Rexec generally terminates when the remote command completes.

You cannot use Rexec to run most interactive commands. For example, you cannot use Rexec to run vi or emacs commands. Use Telnet to run interactive commands.

### Rsh.exe: Rsh

The Rsh connectivity tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Rsh connectivity tool to run commands on remote computers that are running the Rsh service, the UNIX remote shell service. For information about the .rhosts file used to enable this tool, see the description of the Rcp tool earlier in this section.

If the user is logged on to a Windows Server 2003 domain, the domain controller must be available to resolve the user name, because it is not cached on the local computer. Because the user name is required as part of the Rsh protocol, the command fails if the user name cannot be obtained.

## Windows 2003 TCP/IP Technical Reference (Microsoft Corporation)

Rsh copies standard input to the remote command, standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. Rsh generally terminates when the remote command completes.

### **Telnet.exe: Telnet**

The Telnet connectivity tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

The Telnet connectivity tool starts terminal emulation with a remote host running a Telnet server service. The remote computer must also be using TCP/IP. Telnet provides DEC VT 100, DEC VT 52, or ANSI emulation, using the connection-based services of TCP.

To provide terminal emulation from a computer running Windows, the remote host must be running TCP/IP and a Telnet Server service. The Windows-based Telnet user must also have a user account on the remote Telnet Server.

**Note:** Windows Server 2003 and Microsoft Windows XP Professional provide the Telnet Client and Telnet Server components. These are built in, but you must use the Services snap-in to start the Telnet service before it can serve Telnet clients.

### **Tftp.exe: Tftp**

The Tftp connectivity tool is installed with the operating system. This tool is included with all versions of Windows that include TCP/IP.

Use the Tftp connectivity tool to transfer files over the Internet to and from a remote computer (typically, a computer running UNIX) that is running the Trivial File Transfer Protocol (TFTP) service. This tool is similar to Ftp, but it does not provide user authentication. However, the files you transfer using Tftp do require UNIX read and write permissions. You can use Tftp only for unidirectional transfer of files.