



Updates to *Understanding IPv6*

Microsoft Corporation

Published: September 2003

Updated: February 2006

Abstract

This white paper contains updates for changes in Internet Protocol version 6 (IPv6) standards and their implementation in Microsoft® Windows® XP and the Windows Server™ 2003 family, as described in the Microsoft Press® book titled *Understanding IPv6*, by Joseph Davies (ISBN 0-7356-1245-5).

This is a preliminary document and may be changed substantially prior to final commercial release of the software described herein.

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

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

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

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

© 2003 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Microsoft Press, Vista, Windows, the Windows logo, Windows Mobile, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

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

Contents

Global Changes	1
Chapter 1: Introduction to IPv6	2
Chapter 2: IPv6 in the Windows Server 2003 Family	3
Chapter 3: IPv6 Addressing.....	5
Global Addresses.....	5
Site-Local Addresses	6
Zone IDs for Local-Use Addresses	7
Unique Local IPv6 Unicast Addresses.....	8
IPv6 Multicast Scope Definitions.....	8
Chapter 4: The IPv6 Header	10
Chapter 9: IPv6 and Name Resolution	11
Configuration of DNS Traffic Over IPv6.....	11
DNS Traffic over IPv6 Using Locally Configured Unicast Addresses	11
DNS Traffic over IPv6 Using Well-Known Unicast Addresses	11
Source and Destination Address Selection	12
Chapter 11: IPv6 Coexistence and Migration	13
Chapter 12: IPv6 Mobility	14
The Microsoft Mobile IPv6 Technology Preview.....	14
Advanced Networking Pack for Windows XP.....	15
IPv6 Internet Connection Firewall (ICF).....	15
Teredo.....	16
Teredo Components.....	17
Teredo Addresses	18
How Teredo Works.....	19
Changes to IPv6 in Windows XP SP2.....	22
Changes to IPv6 in Windows Server 2003 Service Pack 1	24
IPv6 in Windows Vista and Windows Server "Longhorn"	25
Summary	26
Related Links	27

Global Changes

Throughout the book, the term *Windows .NET Server 2003 family* should be replaced with the term *Windows Server 2003 family*. Microsoft changed the product name after the book was finalized for printing.

Chapter 1: Introduction to IPv6

The following are updates to the information in Chapter 1:

- New domain for IPv6 reverse name space

On page 11, Table 1-1 states that IPv6 uses pointer records in the IP6.INT DNS domains for reverse name resolution for IPv6 addresses. According to RFC 3152, Internet Engineering Task Force (IETF) consensus has been reached that the IP6.ARPA domain be used, instead of IP6.INT. Additionally, RFC 4159 formally deprecates the use of IP6.INT. The IP6.ARPA domain is the domain used by IPv6 for Windows Server 2003. The instance of *IP6.INT* in Table 1-2 should be replaced with *IP6.ARPA*.

Chapter 2: IPv6 in the Windows Server 2003 Family

The following are updates to the information in Chapter 2:

- Internet Protocol security (IPsec) for IPv6 supports the SHA1 HMAC

IPsec for IPv6 in Windows Server 2003 also supports the Secure Hash Algorithm 1 (SHA1) hashed message authentication code (HMAC). To specify the use of SHA1, use the "HMAC-SHA1" string in the .sad file when using the Ipsec6.exe command to specify the entries in the security association database.

- File and printer sharing component can use global addresses

The file and printer sharing components of Windows Server 2003, the Client for Microsoft Networks and File and Printer Sharing for Microsoft Networks services installed in Network Connections, now supports the use of global addresses. Page 25 states that file and printer sharing are only available within a site. This should be changed to state that file and print sharing could use both site-local and global addresses, as long as the global addresses for the client and server computer are from the same site prefix.

- IPv6 support for simple TCP services

Windows Server 2003 includes IPv6 support for both the client and server-side of the simple TCP/IP services. The client-side simple TCP services are available using the Telnet.exe tool. The server-side is installed as Simple TCP Services, an optionally installed Windows component (under Networking Services) using the Add or Remove Programs item in Control Panel.

- IPv6 support for DCOM

Windows Server 2003 includes support for the Distributed Component Object Model (DCOM) application programming interface. DCOM extends the Component Object Model (COM) to support communication among objects on different computers—on a LAN, a WAN, or even the Internet. With DCOM, your application can be distributed at locations that make the most sense to your customer and to the application.

- Temporary addresses disabled by default for Windows Server 2003

The creation and use of temporary addresses, global addresses with a randomly derived interface ID to provide a level of anonymity when communicating over the IPv6 Internet, are disabled by default for Windows Server 2003. Temporary addresses are designed to primarily provide anonymity for client computers and client applications. To enable temporary addresses on a computer running Windows Server 2003, use the **netsh interface ipv6 set privacy state=enabled**.

- Additional Windows XP Service Pack 1 (SP1) and Windows XP Service Pack 2 (SP2) feature set differences

Windows XP SP1 and Windows XP SP2 do not include IPv6 support for file and print sharing (both as a file sharing client and as a file sharing server) and the IPHelper and DCOM APIs.

- DNS traffic over IPv6 disabled by default for DNS Server service

By default, the Windows Server 2003 DNS Server service does not listen for DNS traffic sent over IPv6. To enable the DNS Server service to use DNS over IPv6, use the **dnscmd /config /EnableIPv6 1** command, and then restart the DNS Server service. Dnscmd.exe is part of the Windows Support Tools installed from the \Support\Tools folder on the Windows Server 2003 product CD.

- Basic firewall support for IPv6 traffic

IPv6 for Windows Server 2003 includes support for a basic firewall on an interface. When enabled, IPv6 drops incoming TCP Synchronize (SYN) segments and drops all incoming unsolicited UDP messages. This functionality is disabled by default on all interfaces and can be enabled with the **netsh interface ipv6 set interface interface=*NameOrIndex* firewall=enabled** command. The basic firewall is replaced with Windows Firewall in Windows Server 2003 Service Pack 1.

- IPv6 support for Ttcp.exe

Windows Server 2003 includes IPv6 support for Ttcp.exe, a TCP and UDP diagnostics and troubleshooting tool. Ttcp.exe is available from the Valueadd\Msft\Net\Tools folder on the Windows Server 2003 product CD.

- .NET Framework version 1.1 support for IPv6 is not enabled by default

Although the .NET Framework version 1.1 supports IPv6, it is not enabled by default. To configure the .NET Framework to support IPv6, change **<ipv6 enabled="false"/>** to **<ipv6 enabled="true"/>** in the **<system.net>** section of the Machine.config file in the *Systemroot\Microsoft.net\Framework\v.1.1.4322\Config* folder.

The .NET Framework version 2.0 and later has IPv6 support enabled by default.

Chapter 3: IPv6 Addressing

The discussions in Chapter 3 are based on RFC 2373. RFC 3513 is the new version of the "Internet Protocol Version 6 (IPv6) Addressing Architecture" RFC, which makes RFC 2373 obsolete. RFC 3513 includes new definitions of global addresses, site-local addresses, and scopes for IPv6 multicast addresses.

Global Addresses

The definition of global addresses has changed (pages 52 through 54). The term *aggregateable global unicast addresses* has been changed to *global unicast addresses*. Global unicast addresses are not confined to the prefix 2000::/3.

Figure 1 shows the structure of global unicast addresses currently being allocated by IANA, as defined in RFC 3587.




Figure 1 The global unicast address as defined in RFC 3587

The fields in the global unicast address are:

Fixed portion set to 001 – The three high-order bits are set to 001.

Global Routing Prefix – Indicates the global routing prefix for a specific organization's site. The combination of the three fixed bits and the 45-bit Global Routing Prefix is used to create a 48-bit site prefix, which is assigned to an individual site of an organization. Once assigned, routers on the IPv6 Internet forward IPv6 traffic matching the 48-bit prefix to the routers of the organization's site.

Subnet ID – The Subnet ID is used within an organization's site to identify subnets. The size of this field is 16 bits. The organization's site can use these 16 bits within its site to create 65,536 subnets or multiple levels of addressing hierarchy and an efficient routing infrastructure.

Interface ID – Indicates the interface on a specific subnet within the site. The size of this field is 64 bits.

The fields within the global unicast address create a three-level structure shown in Figure 2.




Figure 2 The three-level structure of the global unicast address

In the new definition of the global address as defined in RFC 3587, there are no longer TLA ID, NLA ID, or SLA ID fields. Therefore, the discussion of subnetting an NLA ID on pages 64 through 69 no longer applies, and the discussion of the subnetting of SLA IDs/Subnet IDs on pages 69 through 73 only applies to subnetting the Subnet ID field of global addresses.

Site-Local Addresses

The definition of the site-local address has changed (page 55). The 38 bits that were set to 0 after the first 10 fixed bits of 1111 1110 11 are now available for use for subnetting within the site. The result is a 54-bit Subnet ID field, as shown in Figure 3.




Figure 3 The new definition of a site-local address in RFC 3513

Because the global address and the site-local address no longer share the same number of bits in the same location, the discussion of using the same subnet identifier for both types of addresses on page 56 is no longer valid unless you are only using the last 16 bits in the Subnet ID field of the site-local address space as the subnet identifier in your organization. Additionally, the discussion of subnetting SLA IDs/Subnet IDs on pages 69 through 73 only applies to the subnetting the Subnet ID field for global addresses and for site-local address schemes that use only the last 16 bits in the Subnet ID field for the subnet identifier.

For Table 3-8 titled "IPv4 Addressing Concepts and their IPv6 Equivalents", the right-hand column for the second entry in the table on page 81 should be changed to "Site-local addresses (FEC0::/10)".

In Appendix D on page 400, the answer to question 9 should be changed to:

- When you use the last 16 bits of the site-local address space as the subnet identifier, the global address and site-local address share the same structure beyond the first 48 bits of the address. In global addresses, the Subnet ID field identifies the subnet within an organization. For site-local addresses, the Subnet ID field can perform the same function. Because of this, when you only use the last 16 bits of the site-local address space for the Subnet ID (setting the upper 38 bits to 0), you can create a subnetting infrastructure that is used for both site-local and global unicast addresses.

In Appendix D on page 401, the answer to question 18 for the "Site-local unicast address" entry should be changed to "FEC, FED, FEE, or FEF".

Note Site-local addresses have been formally deprecated in RFC 3879 for future IPv6 implementations. Existing implementations of IPv6 can continue to use site-local addresses until a replacement has been standardized. A new version of the "IP Version 6 Addressing Architecture" standard is now published as an Internet draft (draft-ietf-ipv6-addr-arch-v4-0x.txt) and includes the deprecation of site-local addresses. This new Internet draft of the standard for IPv6 addressing is destined to obsolete RFC 3513.

Zone IDs for Local-Use Addresses

Unlike global addresses, local-use addresses (link-local and site-local addresses) can be reused. Link-local addresses are reused on each link. Site-local addresses can be reused within each site of an organization. Because of this address reuse capability, link-local and site-local addresses are ambiguous. To specify which link on which the destination is located or within which site the destination is located, an additional identifier is needed. This additional identifier is a zone identifier (ID), also known as a scope ID, which identifies a connected portion of a network that has a specified scope. The syntax specified in RFC 4007 for identifying the zone associated with a local-use address is the following:

Address%zone_ID

Address is a local-use unicast IPv6 address and *zone_ID* is an integer value representing the zone. The values of the zone ID are defined relative to the sending host. Therefore, different hosts might determine different zone ID values for the same physical zone. For example, Host A might choose 3 to represent the zone of an attached link and host B might choose 4 to represent the same link.

For Windows-based IPv6 hosts, the zone IDs for link-local and site-local addresses are defined as follows:

- For link-local addresses, the zone ID is typically the interface index of the interface either assigned the address or to be used as the sending interface for a link-local destination. The interface index is an integer starting at 1 that is assigned to IPv6 interfaces, which include a loopback and one or multiple tunnel or LAN interfaces. You can view the list of interface indexes from the display of the **netsh interface ipv6 show interface** command.
- For site-local addresses, the zone ID is the site ID, an integer assigned to the site of an organization. For organizations that do not reuse the site-local address prefix, the site ID is set to 1 by default and does not need to be specified. You can view the site ID from the display of the **netsh interface ipv6 show address level=verbose** command.

The following are examples of using Windows tools and the zone ID:

- **ping fe80::2b0:d0ff:fee9:4143%3** In this case, 3 is the interface index of the interface attached to the link containing the destination address.
- **tracert fec0::f282:2b0:d0ff:fee9:4143%2** In this case, 2 is the site ID of the organization site containing the destination address.

In Windows XP and Windows Server 2003, the Ipconfig.exe tool displays the zone ID of local-use IPv6 addresses. The following is an excerpt from the display of the **ipconfig** command:

Ethernet adapter Local Area Connection:

```
Connection-specific DNS Suffix . . : wcoast.example.com
IP Address. . . . . : 157.60.14.219
Subnet Mask . . . . . : 255.255.255.0
IP Address. . . . . : 3ffe:ffff:2a1c:2:1cc8:ef1d:1dd9:8066
IP Address. . . . . : 3ffe:ffff:2a1c:204:5aff:fe56:f5b
IP Address. . . . . : fe80::204:5aff:fe56:f5b%4
Default Gateway . . . . . : 157.60.14.1
                             fe80::20a:42ff:feb0:5400%4
```

For the link-local addresses in the display of the **ipconfig** command, the zone ID indicates the interface index of the interface either assigned the address (for IP Address) or the interface through which an address is reachable (for Default Gateway).

Unique Local IPv6 Unicast Addresses

Site-local addresses provide a private addressing alternative to using global addresses for intranet traffic. However, because the site-local address prefix can be used to address multiple sites within an organization, a site-local address prefix address can be duplicated. The ambiguity of site-local addresses in an organization adds complexity and difficulty for applications, routers, and network managers. For more information, see section 2 of RFC 3879.

To replace site-local addresses with a new type of address that is private to an organization, yet unique across all of the sites of the organization, RFC 4193 defines Unique Local IPv6 Unicast Addresses, also known as local addresses. Figure 4 show the structure of local addresses.




Figure 4 The Local address defined in RFC 4193

The first 7 bits have the fixed binary value of 1111110. All local addresses have the address prefix FC00::/7. The Local (L) flag is set 1 to indicate a local address. The L flag value set to 0 has not yet been defined. Therefore, local addresses with the L flag set to 1 have the address prefix of FD00::/8. The Global ID identifies a specific site within an organization and is set to a randomly derived 40-bit value. By deriving a random value for the Global ID, an organization can have statistically unique 48-bit prefixes assigned to the sites of their organizations. Additionally, two organizations that use local addresses that merge have a low probability of duplicating a 48-bit local address prefix, minimizing site renumbering. Unlike the Global Routing Prefix in global addresses, you should not assign Global IDs in local address prefixes so that they can be summarized.

The global address and local address share the same structure beyond the first 48 bits of the address. In global addresses, the Subnet ID field identifies the subnet within an organization. For local addresses, the Subnet ID field can perform the same function. Therefore, you can create a subnet numbering scheme that can be used for both local and global unicast addresses.

Local addresses have a global scope but their reachability is defined by routing topology. Organizations will not advertise their local address prefixes outside of their organizations or create DNS AAAA entries with local addresses in the Internet DNS.

IPv6 Multicast Scope Definitions

RFC 3513 includes new definitions for the values of the Scope field for IPv6 multicast addresses. Table 3-3 on page 59 should be updated to list the values in Table 1.

Table 1 Defined Values for the Scope Field in RFC 3513

Value	Scope

0	Reserved
1	Interface-local scope
2	Link-local scope
3	Reserved
4	Admin-local scope
5	Site-local scope
8	Organization-local scope
E	Global scope
F	Reserved

All instances of the term *node-local scope* in Chapter 3 should be replaced with *interface-local scope*.

Chapter 4: The IPv6 Header

Page 101 states that Microsoft implementations of IPv6 do not support jumbograms. The IPv6 protocol in Windows XP SP2 now includes support for incoming jumbograms at the IPv6 layer. However, there is no support for sending jumbograms or for receiving jumbograms by UDP or TCP.

Page 90 states that the use of the Flow Label field has not yet been defined. RFC 3697 now defines the use of the Flow Label field, requirements for IPv6 source nodes, IPv6 routers, and flow state establishment methods.

Chapter 9: IPv6 and Name Resolution

The following are updates to the information in Chapter 9:

- New domain for IPv6 reverse name space

On pages 230, 231, and 233, the IP6.INT DNS domain is stated as the domain for reverse name resolution for IPv6 addresses. The IP6.ARPA domain is now used, instead of IP6.INT. RFC 4159 formally deprecates the use of IP6.INT. All instances of *IP6.INT* should be replaced with *IP6.ARPA*.

- DNS traffic over IPv6 disabled by default for DNS Server service

By default, the Windows Server 2003 DNS Server service does not listen for DNS traffic sent over IPv6. To enable the DNS Server service to use DNS over IPv6, use the **dnscmd /config /EnableIPv6 1** command, and then restart the DNS Server service. Dnscmd.exe is part of the Windows Support Tools installed from the \Support\Tools folder on the Windows Server 2003 product CD.

Configuration of DNS Traffic Over IPv6

You can configure DNS traffic over IPv6 to use either locally configured or well-known unicast addresses of DNS servers.

DNS Traffic over IPv6 Using Locally Configured Unicast Addresses

In this method, DNS traffic sent from DNS clients and DNS servers over IPv6 is sent to a unicast address locally assigned to the DNS server, such as a site-local or global address configured on the DNS server based on the receipt of a Router Advertisement message. This method requires the following steps:

1. On each Windows Server 2003 DNS server, enable the DNS Server service for DNS traffic over IPv6 as previously described.
2. Obtain the site-local or global addresses of each DNS server by using the Ipconfig.exe tool.
3. Configure each DNS client computer with the unicast IPv6 addresses of your DNS servers using the **netsh interface ipv6 add dns interface=NameOrIndex address=IPv6Address index=PreferenceLevel** command (one command for each DNS IPv6 address). This step can be automated by placing these commands in a login or other startup script.

DNS Traffic over IPv6 Using Well-Known Unicast Addresses

In this method, DNS traffic sent from DNS clients and DNS servers over IPv6 is sent to a set of well-known unicast address manually configured on the DNS server. This method requires the following steps:

1. On each Windows Server 2003 DNS server, enable the DNS Server service for DNS traffic over IPv6 as previously described.
2. Designate which well-known unicast addresses are to be assigned to which DNS servers. The three well-known unicast addresses are FEC0:0:0:FFFF::1, FEC0:0:0:FFFF::2, and FEC0:0:0:FFFF::3.
3. On each DNS server, add one or more of the well-known unicast addresses using the **netsh**

interface ipv6 add address interface=*NameOrIndex* address=*IPv6Address*.

4. Add host routes for the well-known unicast addresses to your routing infrastructure so that the DNS servers are reachable. First, you must add host routes for the DNS server addresses to the neighboring routers of the DNS servers. If you are using an IPv6 routing protocol, configure it to propagate host routes to the non-neighboring routers. If you are using static routers, add host routes with the appropriate next-hop and metric information to all the non-neighboring routers.

Source and Destination Address Selection

For a typical IPv4-only host that has a single interface assigned one IPv4 address and resolves names using DNS, the choice of which IPv4 addresses to use as the source and destination when initiating communication is straightforward. The source IPv4 address is the address assigned to the interface of the host. The destination addresses to which connections are attempted are the IPv4 addresses returned in the DNS Name Query Response message.

For a typical IPv6 host that has multiple IPv6 addresses assigned to multiple interfaces and multiple IPv6 addresses are returned in the DNS Name Query Response message, the choice of the source and destination IPv6 address is more complex. The source and destination IPv6 addresses should be matched in scope and purpose. For example, an IPv6 host should not choose a link-local source address when communicating with a global destination address. Additionally, the possible destination address should be sorted by preference.

To provide a standardized method to choose source and destination IPv6 addresses with which to attempt connections, RFC 3484 defines the following required algorithms:

- A source address selection algorithm to choose the best source address to use with a destination address.
- A destination address selection algorithm to sort the list of possible destination addresses in order of preference.

For more information about the source and destination address selection algorithms defined in RFC 3484, see [Source and Destination Address Selection for IPv6](http://www.microsoft.com/technet/community/columns/cableguy/cg0206.mspx) at <http://www.microsoft.com/technet/community/columns/cableguy/cg0206.mspx>.

Chapter 11: IPv6 Coexistence and Migration

The following are updates to the information in Chapter 11:

- New domain for IPv6 reverse name space

On pages 268 and 299, the IP6.INT DNS domain is stated as the domain for reverse name resolution for IPv6 addresses. The IP6.ARPA domain is now used, instead of IP6.INT. RFC 4159 formally deprecates the use of IP6.INT. All instances of *IP6.INT* should be replaced with *IP6.ARPA*.

- Windows Server 2003 Service Pack 1, Windows XP SP2, and Windows XP SP1 with the Advanced Networking Pack for Windows XP support the Teredo IPv6 transition technology. For more information, see "Advanced Networking Pack for Windows XP" and "Changes to IPv6 in Windows XP SP2" in this white paper.

The Next Generation TCP/IP stack in Windows Vista™ (now in beta testing) and Windows Server "Longhorn" (now in beta testing) is a new implementation of the TCP/IP protocol suite that includes both IPv4 and IPv6 in a dual IP layer architecture as displayed by Figure 11-1 on page 265. For more information, see [Next Generation TCP/IP Stack in Windows Vista and Windows Server "Longhorn"](http://www.microsoft.com/technet/community/columns/cableguy/cg0905.msp) at <http://www.microsoft.com/technet/community/columns/cableguy/cg0905.msp>.

Chapter 12: IPv6 Mobility

Chapter 12 is a detailed discussion of version 13 of the "Mobility Support in IPv6" Internet draft, which, at the time of the writing of *Understanding IPv6*, was the version that some networking vendors including Microsoft were supporting. Mobile IPv6 is now defined in RFC 3775, "[Mobility Support in IPv6](#)", and RFC 3776, "[Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents](#)".

Although the information in Chapter 12 does not reflect the support for IPv6 mobility as defined in RFCs 3775 and 3776, it does properly describe how the correspondent node provided with the IPv6 protocol for Windows Server 2003 works based on version 13 of the "Mobility Support in IPv6" Internet draft.

For an overview of Mobile IPv6 as defined in RFC 3775, see [Introduction to Mobile IPv6](#) at <http://www.microsoft.com/technet/community/columns/cableguy/cg0904.msp>. For a detailed discussion of how the Mobile IPv6 protocol works as defined in RFC 3775, see the [Understanding Mobile IPv6](#) white paper at <http://www.microsoft.com/downloads/details.aspx?FamilyID=f85dd3f2-802b-4ea3-8148-6cde835c8921&displaylang=en>.

For the latest information about the evolving standards for Mobile IPv6, see the [Mobility for IPv6 \(mip6\) Working Group](#) at <http://www.ietf.org/html.charters/mip6-charter.html>.

The Microsoft Mobile IPv6 Technology Preview

A Mobile IPv6 Technology Preview has been developed by Microsoft Research and supports full correspondent node, mobile node, and home agent functionality for computers running various versions of Windows. However, the Mobile IPv6 Technology Preview is not currently available.

Advanced Networking Pack for Windows XP

The [Advanced Networking Pack for Windows XP](#) is a free Web download that provides a set of platform technologies designed to run on Windows XP with SP1 to enable the use and deployment of distributed, peer-to-peer applications based on Internet standards. The update includes a new version of the IPv6 stack, including support for NAT traversal for IPv6 applications. An IPv6 firewall is included to protect the end-user's machine from unsolicited incoming IPv6 traffic, while the peer-to-peer platform makes it simple to write distributed solutions.

Windows Peer-to-Peer Networking is a developer platform that you can use to create peer-to-peer applications for computers running Windows XP with SP1. The Advanced Networking Pack for Windows XP provides the components to run Windows Peer-to-Peer Networking applications. For more information, see [Windows Peer-to-Peer Networking Web page](#) at <http://www.microsoft.com/p2p>. The Windows Peer-to-Peer Networking platform runs exclusively over IPv6.

The enhancements to IPv6 included in the Advanced Networking Pack for Windows XP are the following:

- IPv6 Internet Connection Firewall (ICF)

Computers using Windows Peer-to-Peer Networking must be protected from malicious users that are using IPv6 traffic in the same way that ICF in Windows XP protects computers from malicious users that are using IPv4 traffic.

- Teredo

When connecting to the Internet, many computers running Windows XP are behind network address translators (NATs), which translate traffic between private and public IPv4 addresses. Teredo is a NAT traversal technology that provides unicast IPv6 connectivity across the IPv4 Internet.

All of the functionality in the Advanced Networking Pack for Windows XP is either included or enhanced in Windows XP SP2. For more information, see "Changes to IPv6 in Windows XP SP2" in this white paper.

IPv6 Internet Connection Firewall (ICF)

A firewall provides security by creating a protective boundary between a computer or network and the Internet. You can use IPv6 ICF to dynamically set restrictions for traffic allowed from the Internet. IPv6 ICF is different than the existing ICF in Windows XP, which is used for IPv4 traffic. For more information about ICF for IPv4, see Windows XP Help and Support.

IPv6 Internet Connection Firewall (ICF):

- Runs automatically and provides filters for all network connections on which IPv6 is enabled.
- Monitors all outbound traffic and dynamically creates inbound packet filters for response traffic. This is known as stateful filtering. IPv6 ICF silently discards all unsolicited inbound traffic.
- Logs IPv6 traffic events to a separate log file (from IPv4 ICF). By default, this log file is `systemroot\firewall-v6.log`.

IPv6 ICF protects your computer by silently dropping all packets that are initiated from a source outside the IPv6 ICF computer, such as the Internet. IPv6 ICF stops common attempts to illegally gain access to your computer or network by malicious Internet users using techniques such as port scanning. Instead of notifying you of inbound discarded traffic, IPv6 ICF creates entries in a security log from which you can view the activity that is tracked by the firewall.

IPv6 ICF is configured using commands in the **netsh firewall** context. You can use netsh commands to configure IPv6 ICF to allow specific types of ICMPv6 traffic or traffic to specific TCP or UDP ports. For more information, see [To configure IPv6 Internet Connection Firewall](#).

Note Windows XP only shows the IPv4 ICF configuration in Network Connections. IPv6 ICF may appear disabled, but it is actually enabled and filtering IPv6 traffic.

IPv6 ICF has been replaced with Windows Firewall in Windows XP Service Pack 2. For more information, see "Changes to IPv6 in Windows XP SP2" in this article.

Teredo

Teredo, also known as IPv4 NAT traversal for IPv6, is an IPv6/IPv4 transition technology. Teredo provides address assignment and host-to-host automatic tunneling for unicast IPv6 connectivity across the IPv4 Internet when IPv6/IPv4 hosts are located behind one or multiple IPv4 NATs. To traverse IPv4 NATs, IPv6 packets are sent as IPv4-based User Datagram Protocol (UDP) messages. For more information about how network address translation works, see [Windows 2000 Network Address Translator \(NAT\)](#).

6to4 provides the same function as Teredo; however, 6to4 router support is required in the edge device that is connected to the Internet. 6to4 router functionality is not widely supported by IPv4 NATs. Even if the NAT were 6to4-enabled, 6to4 would still not work for configurations in which there are multiple NATs between a site and the Internet.

Note Windows XP with Internet Connection Sharing and the IPv6 protocol supports 6to4 router functionality.

Teredo resolves the issues of the lack of 6to4 functionality in modern-day NATs or multi-layered NAT configurations by tunneling IPv6 packets between the hosts within the sites. In contrast, 6to4 uses tunneling from the edge device. Tunneling from the hosts presents another issue for NATs: IPv4-encapsulated IPv6 packets are sent with the Protocol field in the IPv4 header set to 41. Most NATs only translate TCP or UDP traffic and must either be manually configured to translate other protocols or have an installed NAT editor that handles the translation. Because Protocol 41 translation is not a common feature of NATs, IPv4-encapsulated IPv6 traffic will not flow through typical NATs. Therefore, the IPv6 packet is encapsulated as an IPv4 UDP message, containing both IPv4 and UDP headers. UDP messages can be translated by most NATs and can traverse multiple layers of NATs.

It is important to note that Teredo is designed as a last resort transition technology for IPv6 connectivity. If native IPv6, 6to4, or Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) connectivity is present between communicating nodes, Teredo is not used. As more IPv4 NATs are upgraded to support 6to4 and IPv6 connectivity become ubiquitous, Teredo will be used less and less, until eventually it is not used at all.

Note Teredo in Windows XP and Windows Server 2003 works only over cone and restricted NATs. A cone NAT stores a mapping between an internal (private) address and port number and an external (public) address and port number. After the NAT translation table entry is in place, inbound traffic to the external address and port number is allowed from any source address and port number.

A restricted NAT stores a mapping between an internal address and port number and an external address and port number, for either specific external addresses or specific external addresses and port numbers. An inbound packet that does not match a NAT translation table entry for both the external destination address and port number and a specific external address or port number is silently discarded. There is an additional type of NAT, known as a symmetric NAT, which maps the same internal address and port number to different external addresses and ports, depending on the external destination address (for outbound traffic). Teredo in Windows XP and Windows Server 2003 does not work over symmetric NATs.

Teredo Components

The set of components that enables Teredo connectivity is shown in Figure 5.

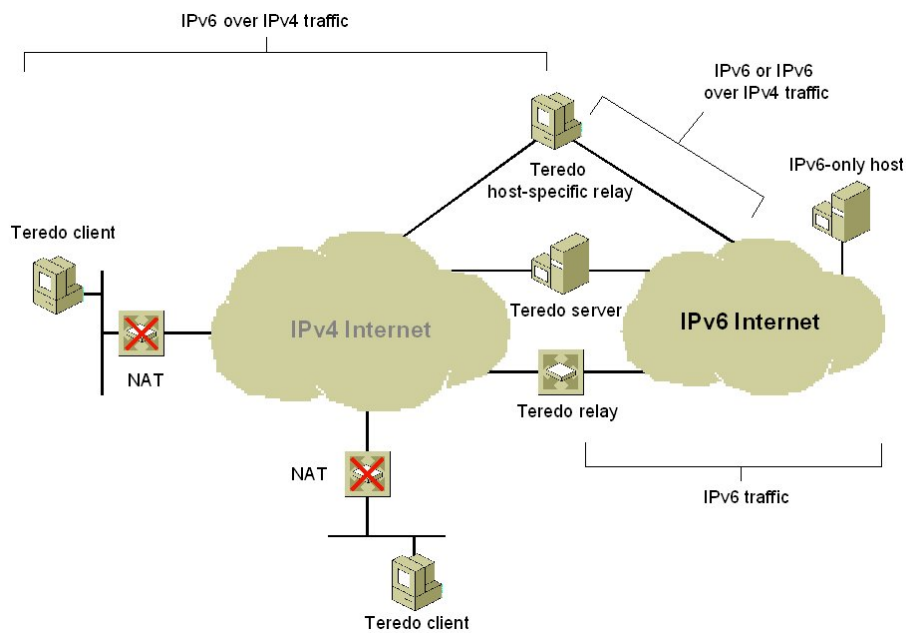


Figure 5 Components of Teredo connectivity

- Teredo client

An IPv6/IPv4 node that supports a Teredo tunneling interface through which packets are tunneled to either other Teredo clients or nodes on the IPv6 Internet (through a Teredo relay).
- Teredo server

An IPv6/IPv4 node that is connected to both the IPv4 Internet and the IPv6 Internet. The role of the Teredo server is to assist in the initial configuration of Teredo clients and to facilitate the initial communication between either different Teredo clients or between Teredo clients and IPv6-only hosts.
- Teredo relay

An IPv6/IPv4 router that can forward packets between Teredo clients on the IPv4 Internet and IPv6-only hosts.
- Teredo host-specific relay

An IPv6/IPv4 node that has an interface and connectivity to both the IPv4 Internet and the IPv6 Internet and can communicate directly with Teredo clients over the IPv4 Internet, without the need for an intermediate Teredo relay. The connectivity to the IPv4 Internet can be through a public IPv4 address or through a private IPv4 address and a neighboring NAT. The connectivity to the IPv6 Internet can be through a direct connection to the IPv6 Internet or through an IPv6/IPv4 transition technology such as 6to4.

The Windows XP Advanced Networking Pack for Windows XP includes Teredo client and Teredo host-specific relay functionality. When the Advanced Networking Pack for Windows XP is installed, the Teredo host-specific relay functionality is automatically enabled if a global IPv6 address has been assigned. A global address can be assigned from a Router Advertisement message that is received from a native IPv6 router, an ISATAP router, or a 6to4 router. If there is no global address configured, Teredo client functionality is enabled.

Teredo Addresses

Teredo addresses have the format shown in Figure 6.




Figure 6 The Teredo address format

A Teredo address consists of the following:

- Teredo prefix

The first 32 bits are for the Teredo prefix, which is the same for all Teredo addresses. The Internet Assigned Numbers Authority (IANA) has not yet defined this prefix, although the prefix 3FFE:831F::/32 is used for initial deployment.

- Teredo server IPv4 address

The next 32 bits contain the IPv4 public address of the Teredo server that assisted in the configuration of this Teredo address.

- Flags

The next 16 bits are reserved for Teredo flags. The only defined flag is the high-order bit known as the Cone flag. The Cone flag is set when the NAT connected to the Internet is a cone NAT.

- Obscured external port

The next 16 bits store an obscured version of the external UDP port that corresponds to all Teredo traffic for this Teredo client. When the Teredo client sends its initial packet to a Teredo server, the source UDP port of the packet is mapped by the NAT to a different, external UDP port. All Teredo traffic for the host uses the same external, mapped UDP port.

The external port is obscured by exclusive ORing (XORing) the external port with 0xFFFF. For example, the obscured version of the external port 5000 in hexadecimal format is EC77 (5000 equals 0x1388, and 0x1388 XOR 0xFFFF equals 0xEC77). Obscuring the external port prevents NATs from translating it within the payload of the packets that are being forwarded.

- Obscured external address

The last 32 bits store an obscured version of the external IPv4 address that corresponds to all Teredo traffic for this Teredo client. Just like the external port, when the Teredo client sends its initial packet to a Teredo server, the source IP address of the packet is mapped by the NAT to a different, external address.

The external address is obscured by XORing the external address with 0xFFFFFFFF. For example, the obscured version of the public IPv4 address 131.107.0.1 in colon-hexadecimal format is 7C94:FFFE (131.107.0.1 equals 0x836B0001, and 0x836B0001 XOR 0xFFFFFFFF equals 0x7C94FFFE). Obscuring the external address prevents NATs from translating it within the payload of the packets that are being forwarded.

How Teredo Works

For two Windows-based computers, the most crucial Teredo processes are those used for initial configuration and communication with a peer in a different site.

Initial Configuration

Initial configuration for Teredo clients is accomplished by sending a series of Router Solicitation messages to Teredo servers. The responses are used to derive a Teredo address and determine whether the client is behind a cone, restricted, or symmetric NAT. If the Teredo client is behind a symmetric NAT, then it cannot function. You can see what type of NAT the Teredo client has discovered from the display of the **netsh interface ipv6 show teredo** command.

Based on the received Router Advertisement messages, the Teredo client constructs its Teredo address from the following:

- The first 64 bits are set to the value included in the Prefix Information option of the received router advertisement. The 64-bit prefix advertised by the Teredo server consists of the Teredo prefix (32 bits) and the public IPv4 address of the Teredo server (32 bits).
- The next 16 bits are either 0x8000 (cone NAT) or 0x0 (restricted NAT).
- The next 16 bits are set to the obscured external UDP port number that is included in a special Teredo header in the router advertisement.
- The last 32 bits are set to the obscured external IP address that is included in a special Teredo header in the router advertisement.

Initial Communication Between Two Teredo Clients in Different Sites

The success of initial communication between Teredo clients located in different sites depends on whether those sites are using cone NATs or restricted NATs.

When both Teredo clients are located behind cone NATs, the NAT translation table entries for Teredo traffic for each Teredo client allows traffic from any source IP address or source UDP port. Therefore, a Teredo client in one site can send packets directly to a Teredo client in another site without the use of additional packets to establish NAT translation table entries.

When the Teredo clients are located behind restricted NATs, additional NAT translation table entries must be established before unicast packets can be exchanged. Figure 7 shows the initial communication process between Teredo clients that are located in different sites when both sites are using restricted NATs.

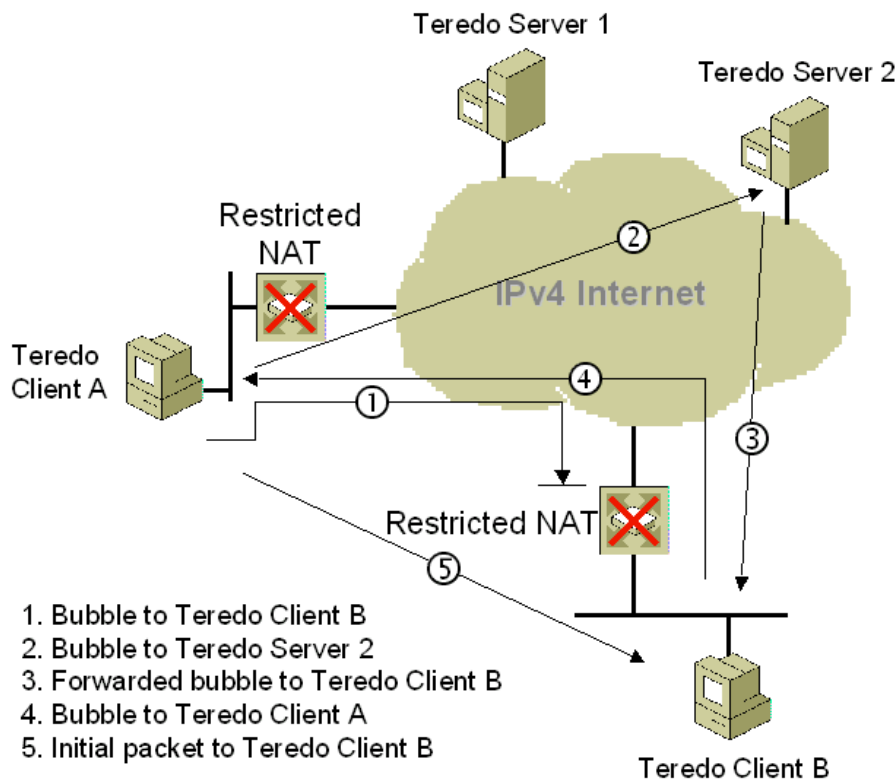


Figure 7 Sending traffic between two Teredo clients behind restricted NATs

To send an initial communication packet from Teredo Client A to Teredo Client B, the following process is used:

1. Teredo Client A sends a bubble packet directly to Teredo Client B. A bubble packet contains no data and is used to create or maintain NAT mappings. Because Teredo Client B is behind a restricted NAT, Teredo traffic from an arbitrary source IPv4 address and UDP port number is not allowed unless there is a source-specific NAT translation table entry. Assuming that there is none, the restricted NAT silently discards the bubble packet. However, when the restricted NAT for Teredo Client A forwarded the bubble packet, it created a source-specific NAT translation table entry that will allow future packets sent from Teredo Client B to be forwarded to Teredo Client A.
2. Teredo Client A sends a bubble packet to Teredo Client B through Teredo Server 2 (Teredo Client B's Teredo server). The IPv4 destination address in the bubble packet is set to the IPv4 address of Teredo Server 2, which Teredo Client A determines from the third and fourth blocks of Teredo Client B's Teredo address.
3. Teredo Server 2 forwards the bubble packet to Teredo Client B. The restricted NAT for Teredo Client B forwards the packet because there is an existing source-specific mapping for Teredo traffic from Teredo Server 2 (established by the initial configuration of Teredo Client B).
4. Teredo Client B responds to the bubble packet received from Teredo Client A with its own bubble packet, which is sent directly to Teredo Client A. Because Teredo Client A's restricted NAT has a source-specific mapping for Teredo traffic from Teredo Client B (as established by the initial bubble packet sent from Teredo Client A in step 1), it forwards the bubble packet to Teredo Client A.

5. Upon receipt of the bubble packet from Teredo Client B, Teredo Client A determines that source-specific NAT mappings exist for both NATs. Teredo Client A sends an initial communication packet directly to Teredo Client B.

This process occurs transparently to the user at Teredo Client A.

There are additional initial communication processes that depend on whether the destination for the initial communication is on the same link, on the IPv6 Internet, or for a Teredo host-specific relay. For more information, see the [Teredo Overview](http://www.microsoft.com/technet/prodtechnol/winxpro/maintain/teredo.msp) white paper at <http://www.microsoft.com/technet/prodtechnol/winxpro/maintain/teredo.msp>.

Changes to IPv6 in Windows XP SP2

Windows XP SP2 includes the IPv6 protocol that was provided with the Advanced Networking Pack for Windows XP, which includes Teredo client and host-specific relay support. Additionally, IPv6 ICF has been replaced with Windows Firewall, a replacement for the Internet Connection Firewall (ICF) component included in Windows XP and Windows XP with SP1.

The Windows Firewall has built-in support for IPv6 traffic and is automatically enabled on all IPv6 connections. Both IPv4 and IPv6 share the same settings for excepted traffic. For example, if you except file and print sharing traffic, then both IPv4 and IPv6-based unsolicited incoming file and print sharing traffic is allowed.

You can configure the Windows Firewall to allow excepted IPv4 and IPv6 traffic in the following ways:

- Windows Firewall Control Panel applet

For more information, see [Manually Configuring Windows Firewall in Windows XP Service Pack 2](#).

- Netsh commands

You can configure Windows Firewall settings through a series of commands in the **netsh firewall** context. Using Netsh, you can create Netsh scripts to automatically configure a set of Windows Firewall settings for both IPv4 and IPv6. For more information, see Appendix B of [Deploying Windows Firewall Settings for Microsoft Windows XP with Service Pack 2](#).

- New configuration APIs

With Windows XP and Windows XP with SP1, there are APIs by which applications can automatically configure excepted traffic and configure ICF settings. With Windows XP SP2, there are new APIs through which you can configure Windows Firewall for global settings for all the items that are available through the Windows Firewall Control Panel applet. You can use these APIs to create customized configuration programs that can be run by users on an organization network. For information about the new Windows Firewall APIs, see [Windows Firewall](#) in the Windows Software Development Kit (SDK).

- Extensive support to configure settings using Group Policy

To centralize the configuration of large numbers of computers in an organization network that use the Active Directory® directory service, Windows Firewall settings for computers running Windows XP with SP2 can be deployed through Computer Configuration Group Policy. A new set of Computer Configuration Group Policy Windows Firewall settings allow a network administrator to configure Windows Firewall operation modes, excepted traffic, and other settings using a Group Policy object. For more information, see [Deploying Windows Firewall Settings for Microsoft Windows XP with Service Pack 2](#).

For more information about Windows XP SP2, see [Service Pack 2 for Windows XP: Resources for IT Professionals](#) at <http://www.microsoft.com/technet/prodtechnol/winxpro/maintain/winxpsp2.msp>.

To protect the computer between when the computer starts and when the Windows Firewall (WF)/Internet Connection Sharing (ICS) service is started, the IPv6 protocol uses a startup firewall policy. Unfortunately, this startup policy blocks incoming Router Advertisement messages, causing a delay in router discovery and address autoconfiguration until the Windows Firewall (WF)/Internet

Connection Sharing (ICS) service is started and the next multicast Router Advertisement is received. The fix for this issue for Windows XP is being considered for Windows XP Service Pack 3. One workaround is to lower the pseudo-periodic interval at which your routers send their router advertisements so that Windows XP SP2-based computers do not have to wait long after the startup policy is removed to perform address autoconfiguration. This issue has been fixed for Windows Firewall in Windows Server 2003 Service Pack 1.

Windows XP SP2 also includes support for incoming jumbograms at the IPv6 layer. However, there is no support for sending jumbograms or for receiving jumbograms by UDP or TCP.

Changes to IPv6 in Windows Server 2003 Service Pack 1

Windows Server 2003 Service Pack 1 includes the Teredo client and host-specific relay and Windows Firewall support for IPv6 traffic that is included in Windows XP SP2.

IPv6 in Windows Vista and Windows Server "Longhorn"

Windows Vista (now in beta testing) and Windows Server "Longhorn" (now in beta testing) include the Next Generation TCP/IP stack, a new implementation of the TCP/IP protocol suite that includes both IPv4 and IPv6 in a dual IP layer architecture.

For information about the features and architecture of the Next Generation TCP/IP stack, see [Next Generation TCP/IP Stack in Windows Vista and Windows Server "Longhorn"](http://www.microsoft.com/technet/community/columns/cableguy/cg0905.mspx) at <http://www.microsoft.com/technet/community/columns/cableguy/cg0905.mspx>.

For more information about IPv6 features in the Next Generation TCP/IP Stack, see [Changes to IPv6 in Windows Vista and Windows Server "Longhorn"](http://www.microsoft.com/technet/community/columns/cableguy/cg1005.mspx) at <http://www.microsoft.com/technet/community/columns/cableguy/cg1005.mspx>.

Summary

This paper details the changes and updates to IPv6 and its implementation in Windows Server 2003 and Windows XP that are not reflected in the Microsoft Press book *Understanding IPv6* by Joseph Davies.

Related Links

See the following resources for further information:

- [Microsoft Windows IPv6 Web site](http://www.microsoft.com/ipv6) at <http://www.microsoft.com/ipv6>
- ["Understanding IPv6" Microsoft Press book](http://www.microsoft.com/MSPress/books/4883.asp) at <http://www.microsoft.com/MSPress/books/4883.asp>
- [Teredo Overview](http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/teredo.mspix) white paper at <http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/teredo.mspix>
- [Understanding Mobile IPv6](http://www.microsoft.com/downloads/details.aspx?FamilyID=f85dd3f2-802b-4ea3-8148-6cde835c8921&displaylang=en) white paper at <http://www.microsoft.com/downloads/details.aspx?FamilyID=f85dd3f2-802b-4ea3-8148-6cde835c8921&displaylang=en>

For the latest information about Windows Server, see the [Windows Server 2003 Web site](http://www.microsoft.com/windowsserver2003) at <http://www.microsoft.com/windowsserver2003>.